# ERCT: An Ontology for Describing Randomised Controlled Trials in the Social Sciences

**Matt Murtagh-White**

**A Dissertation**

Presented to the University of Dublin, Trinity College

in partial fulfilment of the requirements for the degree of

**Master of Science in Computer Science (Intelligent Systems)**

Supervisor: Declan O' Sullivan

Assistant Supervisor: PJ Wall

August 2021

# Declaration

I, the undersigned, declare that this work has not previously been submitted as anexercise for a degree at this, or any other University, and that unless otherwise stated, ismy own work.

Matt Murtagh-White

August 30th, 2021

# Permission to Lend and/or Copy

I, the undersigned, agree that Trinity College Library may lend or copy this thesis upon request.

Matt Murtagh-White

August 30th, 2021

# ERCT: An Ontology for Describing Randomised Controlled Trials in the Social Sciences

Matt Murtagh-White, Master of Science in Computer Science

University of Dublin, Trinity College, 2021

Supervisor: Declan O' Sullivan

Assistant Supervisor: PJ Wall

In the past two decades, the use of Randomised Controlled Trials (RCTs) in economics and international development has grown, providing policymakers and researchers with new insight into what interventions work in improving people's welfare. This dissertation designs and evaluates a Web Ontology Language (OWL) based ontology that seeks to better organise the growing library of RCT evidence for systematic review and link them to international country level development indicators. Gathering competency questions via interviews and a review of systematic reviews, the ontology is created using reusable modular design principles. Using data from the American Economic Association Registry of RCTs, the International Initiative for Impact Evaluation's evidence hub and the World Bank, data is uplifted to fit the ontology to evaluate and validate the design. It was found that the ontology is effective at filtering relevant studies for users but cannot provide quantitative treatment results in its current form due to a lack of data. Future work that gathers these data from alternative sources could alleviate this issue.

# Acknowledgments

I would like to thank my supervisor Professor Declan O' Sullivan and co-supervisor Dr. PJ Wall for their support throughout the writing of this dissertation. For someone entirely new to computer science, their belief and support in helping me take an interdisciplinary idea to its completion with patience and enthusiasm has made the entire dissertation process insightful and enjoyable. I would also like to thank my classmates for all their support throughout a strange and difficult year. We did it.

**Matt Murtagh-White**

*University of Dublin, Trinity College*
*August 2021*

# Contents

# List of Tables

# List of Figures

# Chapter 1

# Introduction

In the last 25 years, a revolution in evidence-based approaches to interventions in economics has occurred. Randomised Controlled Trials (RCTs), long used as a method of testing new drugs in the field of medicine, made its way to economics (Bannerjee and Duflo 2017). Introducing the concepts of control groups, treatment groups and treatment effects into the realm of economics has enriched many areas of the discipline, with top economic academic journals, such as the American Economic Review (AER), the Quarterly Journal of Economics (QJE) and the Journal of Political Economy (JPE) noting an increase in representation of RCTs, with RCT studies growing from 0 to 10 between 2000 and 2015 (Banerjee et al. 2016). During this time, development economists and international development experts gained a great deal of knowledge on the effectiveness of interventions such as microloans, savings and cash transfers. This evidence has helped guide the direction of research in international development and economics and prioritise policies that are most effective at combating poverty. In particular, they are critical in determining which programs are those that donors and policymakers should scale based on their experimental evidence, the vast majority of which has favoured programs which are "low-cost, well-defined, and simple" (Bannerjee & Duflo 2017).

Concurrently, the adoption of experimental methods has allowed the field of behavioural economics to thrive, with researchers in the area achieving an understanding of economic phenomena such as anchoring and mental accounting (Furnham & Boo 2011, Reisch & Zhao 2017). These evidence-based studies are growing exponentially in number; the American Economics Review Registry of RCTs alone now contains more than 4000 studies registered since its inception in 2012.

Yet, with the growing variety and number of studies comes an additional challenge; synthesising this data into a robust meta-narrative that can explain the trends and effects of interventions becomes more and more difficult. Furthermore, recent attention has been focused on the need for replication and the synthesis of evidence in development economics (Sukthankar 2017), the fact that RCTs have not influenced policy as much as expected (Bannerjee et al. 2016), and the need for meta-analysis to synthesise results and guide future research and policy (Duvendack et al. 2012).

A potential answer to modelling the growing data requirements of this field are knowledge graphs. Knowledge graphs are a knowledge representation medium that defines entities through a set of interlinked descriptions defined through a schema which cover a range of topical domains (Paulheim 2017). The concept of the knowledge graph is often linked to the semantic web, which proposes a framework for graph-like data which is presented as objects and subjects, where a predicate describes the relationship between an object and a subject (Berners-Lee 2002). This is expressed as in the Resource Description Framework (RDF), a formal language used to represent subjects as having a predicate or property, the value of which is specified by an object (Manola et al. 2004). As such, each data value in RDF is a triple, with each triple containing an object, predicate and subject. For example, if a person identified as an employee (subject) whose name (predicate) is John Example (object) and has a profession (predicate) of data science (object), then this would represented in an RDF graph as a single triple in figure 1 as follows:


*Figure 1: A graphical representation of the RDF structure*

Data stored in the RDF format can be queried using the SQL-like language SPARQL (Seaborne et al. 2008), which provides functionality to obtain and filter data via object-predicate relationships, among other properties. Many RDF databases, or triplestores, are made available to applications and users via publicly accessible SPARQL endpoints over the internet.

Given the complexity of knowledge structures that can be defined in a graph format, the RDF standard has been extended significantly since its inception. RDF Schema (RDFS) adds subclasses and sub-properties to the RDF format allowing for more complex data modelling based on hierarchies (Brickley et al. 2014). The logical relationship between these objects is further enriched by the Web Ontology Language (OWL), which allows for further specification of groups and checks for consistency (McGuiness and Harmelen 2004). The primary benefit of structuring data in this way is to allow for data from heterogenous sources to be linked, allowing a potentially more complete answer to queries than would otherwise be possible (Berners-Lee et. al 2008, 2011). The expansion of the technology has had a number of successes, one of the most notable being the integration of knowledge graphs into the Google search engine in 2012, which gathers data from a variety of linked data sources in order to deliver collected facts to the user (Ehrlinger & Wöß 2016). The introduction of these technologies has led experts from domains as diverse as history (Shimizu et al. 2020) and medicine (El-Sappagh et al. 2018) to collaborate with knowledge engineers to create OWL ontologies that cohesively describe their subject knowledge.

Alongside the developments in the RDF standard, developing efficient ways of converting traditional tabular data into graph RDF format has also been identified as a priority. Uplifting, the process of capturing and converting variables and relations from tables into a graph format, has had a number of successes: tools such as R2RML (Das et al. 2012) provide a script-based solution to define how data in rows and columns are described in graph form.

In the field of medicine, research has pushed ahead on representing the outcomes of RCTs as knowledge graphs, often defined within an OWL Ontology. This includes ontologies such as the Cochrane databases PICO ontology that represents the populations, methods and outcomes of RCTs (Mavergames et al. 2013), as well as richer, commercially available ontologies and knowledge networks that combine RCT outcomes and biomedical concepts from as much as 200 knowledge sources (Euretos). These innovations have driven research in the area, with semantic properties of these knowledge graphs showing promise in aeras such as prioritising drugs that can be repurposed for treating diseases (Malas 2019). While these advances have been made in medicine, to the authors knowledge, no corresponding ontologies or subsequent related research has been developed for RCTs in international development or economics. The research presented in this paper therefore identifies and corrects this research gap by developing the first ontology that can describe the data and characteristics of RCTs in this field, providing a starting point for the development of semantic applications and research.

## 1.1 Motivation

The motivation for this research is to provide data that can be used to inform interventions in a wide variety of contexts, with particular focus on the global south. Given that the vast majority of RCTs are concentrated in South Asia, East Africa, South and Central America and Southeast Asia (Cameron et al. 2016), the data gathered in this area provides us with a great deal of insight into the lives of those in poverty. Perhaps more importantly, understanding the dynamics of this insight has the potential to guide interventions such that they are most effective in alleviating poverty and to provide a higher quality of living for the worst off. The effects of understanding these dynamics have already been felt, be that by designing a deworming program to best improve both health and educational outcomes (Miguel & Kremer 2004) or providing cash transfers with conditions that can nudge individuals towards more positive educational and employment outcomes (Rawlings & Rubio 2005). By providing a platform to access all records of RCTs and evaluation studies in the contexts required by the researcher, it is hoped that this research can contribute towards further understanding of the available evidence.

## 1.2 Aims

The two main aims of this research are as follows:

a) Design an OWL ontology that describes the methodology, concepts, and results of RCTs in the field of economics

This aim has potential research relevance to the fields of semantic web research, international development and economics. In semantic web research, this ontology will act as a foundation for applications that can assist social scientists in the systematic review process. In international development and economics, through these downstream applications, this work has the potential to assist researchers in synthesising existing research and discovering new research paths. By allowing researchers in international development to access RCT and program evaluation data across all available data providers and use up-to-date country level development indicators to filter relevant work, the ontology provides opportunities to evaluate and review the areas of research that show the most promise.

Finally, this aim has relevance to policy that stems from evidence-based research. The impact and evaluation evidence gathered from RCTs are often integral in deciding whether or not to scale up development intervention programs to a national or international level (Bannerjee et al 2016), as the treatment effects of the intervention central to RCTs on key variables such as consumption, savings and wellbeing are seen as evidence as to whether or not the intervention is working. Given that observed successes in RCTs may be specific to the context they were performed in, replication of successful studies is encouraged to ensure that the effect observed across many contexts, if not all. By combining this data in a knowledge graph, there is potential to discover and scale the interventions that are most effective for a given problem at a rate faster than currently possible. Thus, decision makers could make more evidence backed decisions using a deeper and more relevant pool of evidence from the literature.

OWL was chosen to represent this ontology for several reasons. Foremost of these is that the ontology language provides a number of additional conditions that can be better used to describe the relationship between objects over RDFS. With OWL, cardinality and existence constraints can be expressed between objects, meaning that it is possible for an object to require another object of a given type. There are characteristics of RCTs that require this axiom; for example, RCTs should be conducted by at least one researcher, should have just one methodology and require zero or more outcomes. Similarly, OWL allows for localised range and domain constraints; RCT titles should be titles when a hasName of RCT, while country names should be country names when hasName is applied to countries.

b) Validate the ontology by uplifting existing datasets to match the objects, predicates, and subjects of the newly designed ontology and evaluating against competency questions

The second aim will be to validate the ontology design by testing it within an evaluation framework, determining whether the design can sufficiently resolve the defined modelling issues. The most common best-practice approach to this is by competency questions, a set of questions derived from modelling issues that are used as unit tests for a developed ontology (Tudorache 2020). In order to achieve this validation, a triplestore must be populated with data that conforms to the ontology.

Existing databases exist that contain information on the results, outcomes and methodologies of RCTs in economics, social sciences and international development are available but do not provide an ontological foundation for semantic web development (American Economic Association 2021, International Initiative for Impact Evaluation 2021). Uplifting, the process of converting existing tabular data into an RDF graph format, would allow for these existing databases to be linked and accessible in RDF form. Using this, a sustainable data pipeline can be created to uplift existing databases into the RDF format within the construction of the designed ontology and test their efficacy using a set of competency questions. This will allow for the ontology to be evaluated and lay the groundwork for a functioning SPARQL endpoint, allowing users to query data from each database that conforms to the ontology in a manner that is aligned with their use cases. This endpoint would therefore serve as a platform through which further semantic web applications that use this data could be created.

## 1.3 Research Question

Given the aims and motivation, the research question of this dissertation can be defined as:

***To what extent can the methodology, outcomes and contributions of randomised controlled trials in economics be represented and queried as a knowledge graph?***

## 1.4 Research Objectives

To achieve the defined research aims, this research seeks to fulfil the following objectives:

1. *Gather data and develop ontology competency questions for design and evaluation*

   This objective is key to both the design and evaluation of the ontology. Following best practice (Tudorache 2020), competency questions will be developed through interview of users with domain knowledge in RCTs. Additionally, a review of work in the area of systematic reviews will be conducted to provide further insight into potential use cases. These questions will inform both the design of the ontology as well as the evaluation framework used to validate it.

2. *Uplift existing data to conform to the ontology*

   Existing RCT data from the American Economic Association will be adjusted to fit the completed ontology. In some cases, this involves re-classifying a sample of the data to fit the fixed vocabulary of the ontology. Mapping is completed through the use of R2RML, a framework that allows for tabular data to be uplifted to RDF through a specified mapping (Sundara & Cygniak 2012).

3. *Validate the ontology competency questions*

   Based on the competency questions, a set of SPARQL queries are developed to validate the ontology against its intended design

## 1.5 Data

The data used in this study comes primarily from the AEA RCT Registry and Registry for International Development Impact Evaluations (RIDIE). This data includes descriptions of the methodology, outcomes, locations and topic areas of RCTs, represented as a .csv file. The data requires require some adjustment; outcomes and interventions need to be harmonised to the classifications of the controlled vocabulary. Many of these classifications can be inferred through relatively simple data wrangling, the process of which is outlined further in chapter 4. While further inference of classifications is likely possible from these data sources through the use of more complex classification algorithms, this is outside of the scope of this dissertation and is therefore considered future work in this area.

The secondary source of data used in this study comes from the World Bank and consists of a variety of development indicators used to describe the living conditions and progress of a given jurisdiction (World Bank 2021). This data is obtained through use of an API, more details of which is available in chapter 4.

## 1.6 Expected Contribution

This research will contribute to a number of bodies of literature and have impact for a number of users. Primarily, it will be a foundation point for semantic web research in the domain of RCTs for international development and economics. Downstream research applications that could develop on this ontology could include more efficient information retrieval based on semantic queries, or machine learning applications such as policy recommendations based on graph learning methods such as RDF2Vec (Ristoski et al. 2019). Given that this ontology represents the first step in defining RCTs in economics and the social sciences in OWL, it opens the door to further research on developing ontologies to define a wider range of economic or international development indicators.

A second set of beneficiaries of this research are researchers in economics and international development. The ontology presented in this dissertation is intended to provide linked RCT data from across multiple isolated databases, presented in a structure that prioritises those variables identified as most important by users in the comparison of evidence. The ontology is therefore intended to assist the systematic review process by providing these users with platforms that can quickly and effectively filter relevant study outcomes for inclusion in systematic review.

Finally, this would benefit policymakers by speeding up the process by which they discover which interventions are most effective based on evidence. Evidence from RCTs is viewed as a marker for success, commonly used in deciding whether to scale up intervention programs. This ontology could assist policymakers by allowing for relevant evidence from across these intervention programs to be more easily accessible, allowing evidence based public policy to be more easily formed.

Finally, this research can also benefit providers of RCT data, such as the American Association of Economics or the International Initiative for Impact Evaluation (I3E), in that it would allow them to provide a linked data platform for the first time. This would mean that heterogenous data from both sources could be included in search results for users, ultimately increasing access to RCT evidence and providing them in an amalgamated format that can be immediately analysed. Furthermore, by incorporating the ontology in data provision and search, the providers could potentially link to other isolated databases for more contextual data, such as the World Bank development indicators or macroeconomic data from the International Monetary Fund.

## 1.7 Outline of Dissertation

The dissertation is structured as follows: chapter 2 will review background and related work, outlining where this research contributes to the literature. Chapter 3 discusses the design of the ontology, including a review of the design patterns and design choices employed in creating the ontology. Chapter 4 discusses the implementation of the ontology, detailing the steps taken to produce the data pipeline. Chapter 5 evaluates the completed ontology, outlining a framework of evaluation before applying it to the completed pipeline. Finally, chapter 6 draws conclusions from the research.

# Chapter 2

# Background and Related Work

The literature reviewed for this research is drawn from concepts from four distinct bodies of literature: knowledge and data engineering, bioinformatics, economics, and international development. Each of these areas represent domain and technical knowledge that is relevant to the ontology: knowledge and data engineering provides the technical frameworks that the ontology uses to function, bioinformatics provides examples of existing RCT ontologies and their applications, the ontology is built for data that has been generated from RCTs in economics and the vast majority of these studies are focused on the global south and are relevant to international development. As such, this section will discuss each of their contributions to this research topic in turn.

## 2.1 Knowledge Graphs and Ontology Engineering

From knowledge and data engineering, the chief contributing concepts to this dissertation are knowledge graphs and ontologies. Knowledge graphs evolved from the concept of the semantic web, proposed by Berners-Lee (2002), which outlines a network of data defined by objects, predicates and subjects that describe the data points and their relations to each other. Extensions to this framework, such as RDFS, further enrich the scope of knowledge representation by including subclassing (Brickley et al. 2014). Ontologies, normally defined in the Ontology Web Language (OWL), are used to describe the logical relationships between objects and can be used to validate the integrity of knowledge graphs. The query language SPARQL can be obtain data from such graphs and includes some limited functionality to process data for analytics.

These ontologies have been used to define a wide range of data, from individuals involved in the Atlantic slave trade (Shimizu et al 2020) to social networks (Goldbeck and Rothstein 2008). Shimizu et al. (2020) employ OWL to develop the “Enslaved” ontology, which links a variety of databases focused on documenting individuals involved in the Atlantic slave trade. This ontology was developed by eliciting competency questions from a variety of users of differing skill level, ranging from primary school children to academic historians. Goldbeck and Rothstein’s (2008) work on the “Friend of a Friend” network provides an ontology that can be used to link online social networks, with objects and predicates designed to describe how individuals within this network are related to each other. The latter ontology has resulted in a number of semantic applications: Tang et al. (2008) develop Arnetminer, a network of researchers based on data scraped from a variety of academic databases, while Rowe (2009) uses the ontology to generate insight on user activity across different online social networks.

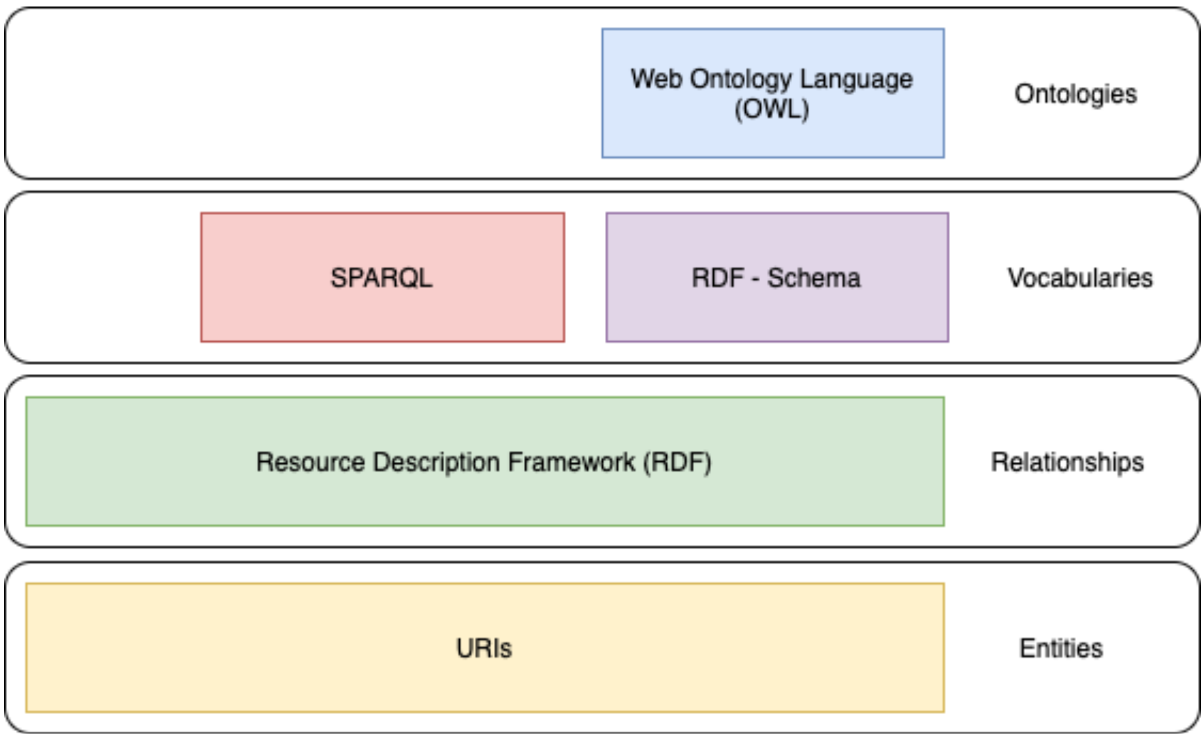

*Figure 2: The semantic web stack*

The development of the OWL standard has been followed by the emergence of research in ontology engineering, where semantic web researchers develop ontologies that that structure and define the axioms and structure knowledge as objects, predicates and subjects in a given domain. Within ontology engineering, the construction of these ontologies has led to the development of best practice principles in the design of ontologies. Gangemi and Presutti 2009 propose a “modular” design philosophy for ontologies, whereby ontologies consist of smaller patterns of objects, predicates and subjects are conceptually packaged as modules that do not have implications for reasoners or queries but allow for a more human-understandable ontology structure. The eXtreme Design Methodology (Presutti et al. 2009) for ontology design furthers these principles by emphasising the reusability of modules across multiple ontologies. These design principles are expounded on and an account of their application is provided in chapter 3.

## 2.2 Structuring Experimental Evidence Within Knowledge Graphs

Researchers have found success in the modelling and structuring the results and methodology of trials in evidence-based medicine as knowledge graphs. The Population, Intervention, Comparison and Outcome (PICO) ontology records a variety of variables on RCTs, and is modelled on the Cochrane Library, a popular library of RCTs in evidence-based medicine (Mavergames et al. 2013). Notably, it includes standardised outcomes, interventions and drugs used in RCTs, among others. Other ontologies in this area are more specialised: the Cancer Care Treatment Outcome Ontology (CCTOO) profiles the treatment outcomes for patients with tumours, consisting of 1133 trials gathered from clinical trials (Lin et al. 2018), while the Ontology of Clinical Research (OCRe) describes the entire lifecycle of RCTs, including their design, execution, reporting, interpretation and application (Sim et al. 2014).

Given the maturity of ontology engineering in this domain, there have also been significant successes in applying novel machine learning methods to these knowledge graphs. One of the premier applications within this domain has been the use of knowledge graphs to suggest potential repurposing of drugs for new diseases. Himmelstein (2017) uses the RDF2Vec tool to transform a knowledge graph consisting of more than 450,000 nodes representing 755 existing treatments into vectors, and using a logistic regression model with elastic-net regularisation, calculates probabilities of treatment for a 209,168 disease compound pairs, using an external study on epilepsy for validation. The results provided significant pharmacological insights on epilepsy, suggesting a number of drugs from clinical trials that were either prematurely stopped or were intended to treat a different disease that might be repurposed for treating the condition.

Similarly, Malas et al. (2019) combine a knowledge graph that integrates over 200 existing biological knowledge bases with the drug repurposing database RepoDB to perform a random forest based analysis that suggests Mozavaptan as a potential candidate to treat Polycystic Kidney Disease. In a similar study, Lei et al. (2020) implement a knowledge graph solution to the classification of noncommunicable diseases, designing a data collection model that incorporates both structured and unstructured medical data into a knowledge graph and vectorising the contents of the graph as in Himmelstein (2017) in order to create features for a machine learning model. Using this vectorised dataset on a variety of classifiers such as Support Vector Machines (SVMs), the authors find that their knowledge graph solution can provide sufficiently accurate insight into the trends of morbidity and mortality through classification.

## 2.3 The Use of Randomised Controlled Trials in Economics

Many areas of economics have benefited greatly from ideas borrowed from evidence-based medicine, maturing from a largely theory-based endeavour to a practice of testing and revising hypotheses in response to evidence gathered from RCTs. Of these fields, none have embraced the practice of RCTs more than development economics. By 1997, large poverty alleviation programs such as PROGRESA were integrating random assignment-based evaluation into their activities, providing a trove of data for development economists and development professionals to test theories against (Banerjee et al. 2016). Economists could now evaluate the effect of interventions on a variety of key economic variables associated with welfare, such as consumption, investment in education and health, household savings and interhousehold bargaining.

From this, international development has learned a variety of lessons that shape policy and research to this day; that childhood exposure to conditional cash transfer schemes can increase educational attainment, mobility and labour market outcomes (Parker and Vogl 2018), that the length of exposure to these programs can improve long term increases in educational attainment (Kugler & Rojas 2018), and that cash obtained in transfers tends to be invested in productive assets that can increase consumption in the long run (Gertler et al 2017) to name just a few. In the years since, the practice of RCTs in development has increased substantially, with focuses turning towards meta-analysis to distil the wide array of data into new directions for policy and future research, usually in the form of domain specific, systematic reviews of the research (Baird et al 2013). Efforts have been made to develop common systematic review frameworks to provide evidence in support of development policies and interventions (Waddington et al. 2012), especially with external validity being a common criticism levied at such research (Peters et al. 2016).

While development economics is by far the most prolific in the use of RCTs, other branches of economics also incorporate them or similar practical experiments into their research. These studies have allowed for the observation of ubiquitous social phenomena such as anchoring, whereby individuals are biased in their choice of value by some previously prevented value (Furnham & Boo 2011), or loss aversion, where individuals are observed to weigh potential losses higher than potential gains (Novemsky and Kahneman 2005). Research in this area contributed to the idea of "nudging" as a way of subtly encouraging positive behaviour in people, a concept popular to policymakers throughout the world (Thaler and Sunstein 2009, Oliver 2013, Hampton & Adams 2018).

Despite the growing practice of systematic review and meta-analysis in the area of international development and economics (Waddington et al. 2012), no effort has yet been made to map the concepts, variables and outcomes of these RCTs to a common ontology. The American Economic Association (AEA) RCT registry contains useful information on these topics but contains little to no standardisation, with many fields in the registry going unused or containing long passages of text that are difficult to parse (American Economic Association 2021). Taking this key step would ensure that meta-analysis and review would take place under a common, comparable structure, and would allow for the development of algorithms and models that can assist with systematic review, recommend new avenues of research and suggest actions for policymakers. Thus, this research will contribute to the literature by developing and evaluating the first ontology in this area, providing an indication of whether this avenue of research might be an effective tool in conducting systematic review.

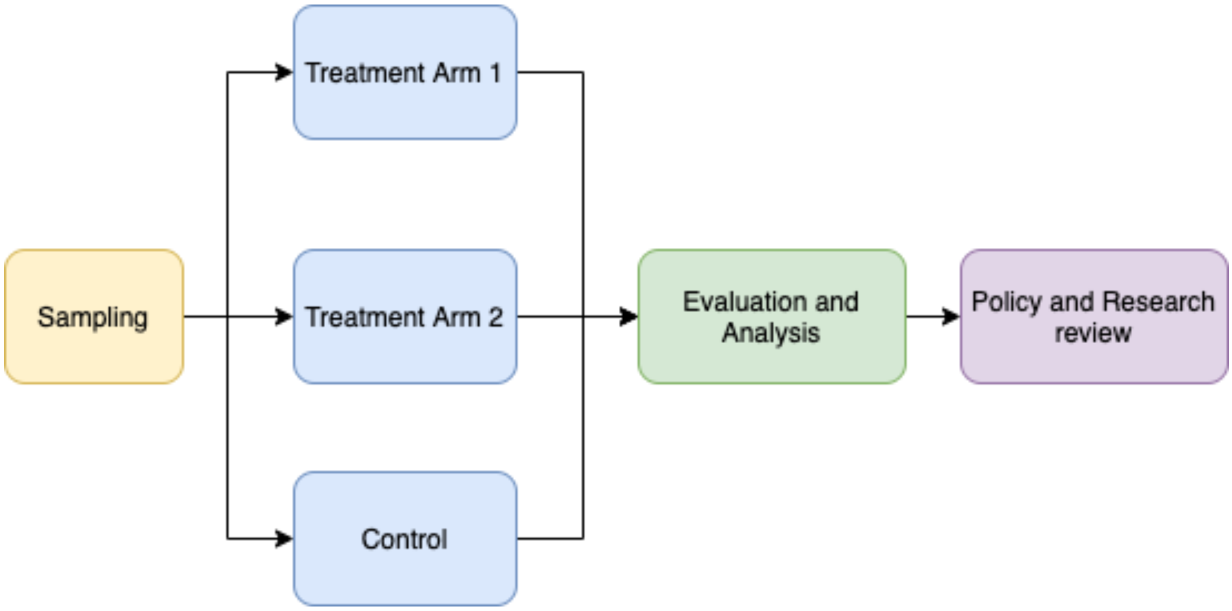


*Figure 3: Typical Lifecycle of an RCT*

## 2.4 Knowledge Graphs in the Global South

While the proliferation of knowledge graphs has led to wide variety of applications in advanced knowledge domains such as medicine or biology, one of the more interesting recent developments in the technology has been adapting it for improving the lives of those living in poverty. These applications of knowledge graphs tend to either focus on increasing the transparency or availability of data commonly used in international development or adapting or applying knowledge graphs in a way that is directly used and benefits those living in poverty.

For the former, a wide array of knowledge graphs has been developed that improve accessibility and transparency in data generated by governments. While many of these, such as data.gov and data.gov.uk, provide access to significant fraction of government data that has been tagged to with RDF metadata (Shadbolt and O'Hara 2013), these initiatives have largely focused on governments in highly developed countries and have not yet seen widespread adoption in the global south. Furthermore, while many international organisations provide extensive data catalogues on countries across the world, most do not yet implement this data on a knowledge graph. For example, research into providing the United Nations Sustainable Development Goals (SDG) indicators in knowledge graph form is only in its infancy and has yet to be offered to the public (United Nations 2020).

Knowledge graph development and adoption in these settings has instead been pushed more from non-governmental organisations. One example of these is the International Aid Transparency Initiative (IATI), an organisation that focuses on the transparency of international development aid. Brandt and de Boer (2015) adapt a standard of publication and aid information developed for the organisation's open data portal into a linked data model. By linking the data to datasets such as the World Bank development indicators and DBPedia, the authors develop a number of applications that use this combined data to demonstrate the ability of the platforms to provide insight on aid transparency.

While these initiatives are likely to have indirect effect on those living in poverty, they are more focused on international organisations and policymakers rather than those living in poverty themselves. This critique of the technology has been propounded by Guéret et al. (2014), who argue that the implementation of knowledge networks and linked data technology has thus far been biased towards the global north; linked data projects have been on a large scale, often combining data from government projects outside the scope of the global south and favouring methods of access that are not common among the global poor. It is argued that to make these technologies relevant to those living in these conditions, they should be "downscaled" to focus on more relevant data such as weather or soil information and be accessible on basic GSM devices as has been the case in the success of mobile money (Jack & Suri 2011).

This initiative has been followed by some success: Valkering et al. (2016), recognising the difficulty in accessing the internet and obtaining access to expensive server and network infrastructure in less developed rural settings develop SPARQL over SMS, which allows for the exchange of RDF data using SMS over a standard GSM network. This allows for access to knowledge graphs without the need of a smartphone or cellular internet. This is the main research focus of de Boer (2021), who focuses on expanding this technology to provide market information to smallhold farmers (De Boer et al 2015) and veterinarian services and information for isolated communities (De Boer 2016), both of which can be accessed through voice-to-text technology to alleviate issues of illiteracy.

Within the scope of the application of knowledge graphs in the global south, the research presented in this paper is focused more towards the former; knowledge graph applications that are relevant for policymakers, international organisations and non-governmental organisations working in the field of development. The research presented here is akin to Brandt and de Boer (2015) in that it concerns the development of a knowledge graph that combines policy-relevant microdata with wider international development indicators such as those provided by the world bank. Finally, it should be acknowledged that the research presented here is not exclusively relevant to policy in the context of the global south – many RCTs, particularly those concerned with behavioural economics, take place in advanced economies and are used in the formation of policy in those countries (Oliver 2013, Hampton & Adams 2018).

# Chapter 3

# Design

The design of the ontology follows the example of Shimizu et al. (2020) in applying the principles of modular ontology design principles derived from established ontology design patterns (Gangemi and Presutti 2009) and eXtreme Design Methodology (Presutti et al. 2009). In short, these design principles involve a modular evolution of the ontology, with each module taking the form of an ontology design pattern. This design direction was chosen due to being the most efficient to scale of those considered; the conceptual disambiguation of modules within the ontology allows for multiple patterns to be incorporated in a way that allows quicker human understanding.

While modules are not differentiated insofar as the organisation of the OWL ontology itself is organised and constituted, a modular organisation of sub-patterns allow the design of the ontology to better capture human conceptualisations of data. This was chosen over other families of patterns such as Gangemi's (2005) Conceptual Ontology Design Patterns (CODeP), which offers a greater degree of generalisation at the cost of an increased implementation complexity. Holsapple and Joshi (2002) make more general recommendations as to how to design an ontology that lacks some of the details that this ontology will require, however, their direction on collaborative design serves as an integral design direction for the user interview step discussed later in this dissertation.

The potential disadvantages of the modular design direction are that the complexity of the ontology increases with the number of modules implemented, and the ontology still faces the trade-off of domain specificity versus usability, which is discussed in more detail later in this section.

Consistent with the chosen approach, this study first employs the steps outlined by Hitzler and Krisnadhi (2018) and later extended by Shimizu et al. (2020). This can be defined as the following:

> *1 Defining use cases and/or scope for the ontology*
>
> *2. Collecting competency questions while looking at possible data sources and scoping the problem*
>
> *3. Identifying structural components in the data that can act as modules, and determining design patterns that can be used for each module*
>
> *4. Join modules together while defining the logical axioms between them*
>
> *5. Create the OWL files in Protégé*

For the first, the use case and scope is largely outlined in the first section of this dissertation and will be only briefly expanded on here. For the second, this section will discuss in detail the process of developing competency questions and how they feed into the evaluation of the ontology. Finally, it will outline the process of developing the key core modules of the ontology as they relate to the fulfilment the use cases and the competency questions.

## 3.1 Analysis of Use Cases and Scope

All ontologies are designed with the intention to structure domain knowledge to allow users to quickly understand and process the underlying data (Hitzler and Krisnadhi 2018). In designing the ontology, the key trade-off in achieving this goal comes with the deciding specificity: should ontological commitments be made to make relations specialised and easily understandable, or should the ontology be designed with broader relations to allow for easier expansion in the future?

In answering this question, the specificity of the domain is first considered. Existing solutions in the areas of RCTs, such as Cochrane's PICO ontology, favour the former approach. This ontology largely formalises a well-established knowledge representation paradigm in evidence-based medicine, which organises data by problem/population, intervention, comparison, and outcome (PICO) (Huang et al. 2006). Even without considering the prevailing criticisms of structuring clinical data within this framework[1], by tightly structuring the data within an existing ontology designed for traditional data, the ontology is limited both in its scope for expansion and the ability to link to other datasets, one of the key features of the RDF/OWL stack (Berners-Lee et. al 2008, 2011). The Ontology of Clinical Research (OCRe) (Sim et al. 2014) corrects this to a certain extent, allowing links with several other medical ontologies contained on the National Centre for Biomedical Ontology's online biomedicine platform, Bioportal (2021). However, despite these developments, an absolute commonly accepted standardisation does not yet appear to have emerged (Sim et al. 2004). Thus, while an ontology such as PICO serves as a useful entry point for researchers familiar with the knowledge format, it sacrifices a certain degree of flexibility for specificity – it would be difficult to see other forms of evidence being incorporated in such an ontology in the future.

On the other hand, ontologies such as that used by DBpedia (2021) are the opposite in scope. In defining data as broadly as possible, the DBpedia ontology allows for almost every data table hosted on Wikipedia to be represented in a queryable form. While this allows for nearly any type of data to be queried, it sacrifices the ability to quickly query specific data, and increases the likelihood of an object being misclassified within the ontology. An ontology such as this likely to be too general for the use case: users are likely to prefer specificity with some limited links to other relative data sources rather than highly linked data at the cost of ease of use.

The question then turns to where an ontology to represent empirical knowledge from RCTs might fall within this range. While PICO serves as a useful starting point, the fundamental difference between ontological modelling of evidence-based medicine and economics is that

[1] Huang et al. find that with a set of 59 clinical competency questions, only two contain all four elements of PICO, while just 37% include both intervention and outcome. Similarly, Schardt et al. found that structuring data within PICO only marginally improved the precision of competency questions compared to PubMed, though the sample size used was quite small.

medicine is already an ontologically rich discipline, where the role of ontologies is consistently debated and discussed (Michie and Johnston 2017). A similarly rich debate is missing in economics and the social sciences in general, for a variety of reasons. The foremost of these, identified in one of the few works produced on the topic, is that the theoretical foundations of social scientists such as economists have, in the past, not been solely determined by empirical evidence (Maki 2001). Rather, the practice of using RCTs as a way to generate data in support of both these theoretical foundations is relatively new (Duflo and Banerjee 2017), with the number of experiments growing at a rate faster than any discussion on how the knowledge output might be structured. Those ontologies that do exist reflect the time when the theoretical foundations of policy and models were less tested and less concerned with the practicalities of experimental testing, such as sampling, experimental design and treatment effects. These range from the JEL codes used by the Journal of Economic literature, which describe the area of study in broad terms, to the thesaurus of economics, an attempt by the Leibniz Information Centre for Economics to establish a hierarchical classification of economic terms (Neubert 2009).

## 3.2 Design Scope and Principles

From this reflection of existing solutions and best practices, the design of this ontology aims to combine four principles to define design scope. In summary:

a. *The ontology will first and foremost be intended to assist with systematic review.*

The primary intention of the ontology is for systematic review; it aims to allow users to discover when and under what conditions interventions are most effective and why, with useful differentiation based on properties such as research area, methodology and geography.

b. *The trial data will be structured on established principles from PICO, widely used in evidence-based medicine.*

The ontology will incorporate elements of the structure of PICO from bioinformatics, with particular emphasis on population, intervention and outcome. Despite the discussed shortcomings, the ontology is still favoured by established knowledge bases of trial data, and its principles are still largely relevant to the RCTs as they are applied to the social sciences.

c. *This structure will be adjusted to incorporate the practicalities and established theoretical foundations and classifications of the social sciences.*

The ontology will be framed around the both the practicalities and established classifications of evidence-based social science, an area yet to be widely explored in an ontology.

d. *The ontology must use modular design patterns to allow these terms to expand without an entire redesign of the ontology.*

To achieve this design principle c, the fourth principle is applied: the ontology must use design patterns that allow these social science terms to expand without requiring an entire redesign of the ontology (Hitzler and Krisnadhi 2018).

e. *Patterns used in modules, wherever possible, are generic solutions that solve localised problems (Presutti et al. 2009).*

Presutti et al. characterise ontologies as having a “problem space”, that are composed of the actual problems posed to users that the ontology attempts to resolve, and “solution space”, made up of reusable modelling solutions and microformats that can be matched to the issues identified in the problem space. In pursuit of this goal, the modules identified within the data are matched to established design patterns, or where none exist, a new design pattern for a module is created with the principle that is may be reused for future ontologies. This approach acknowledges and incorporates more recent criticism over the reusability of ontologies (Fernández-López et al. 2019).

## 3.3 Use Cases

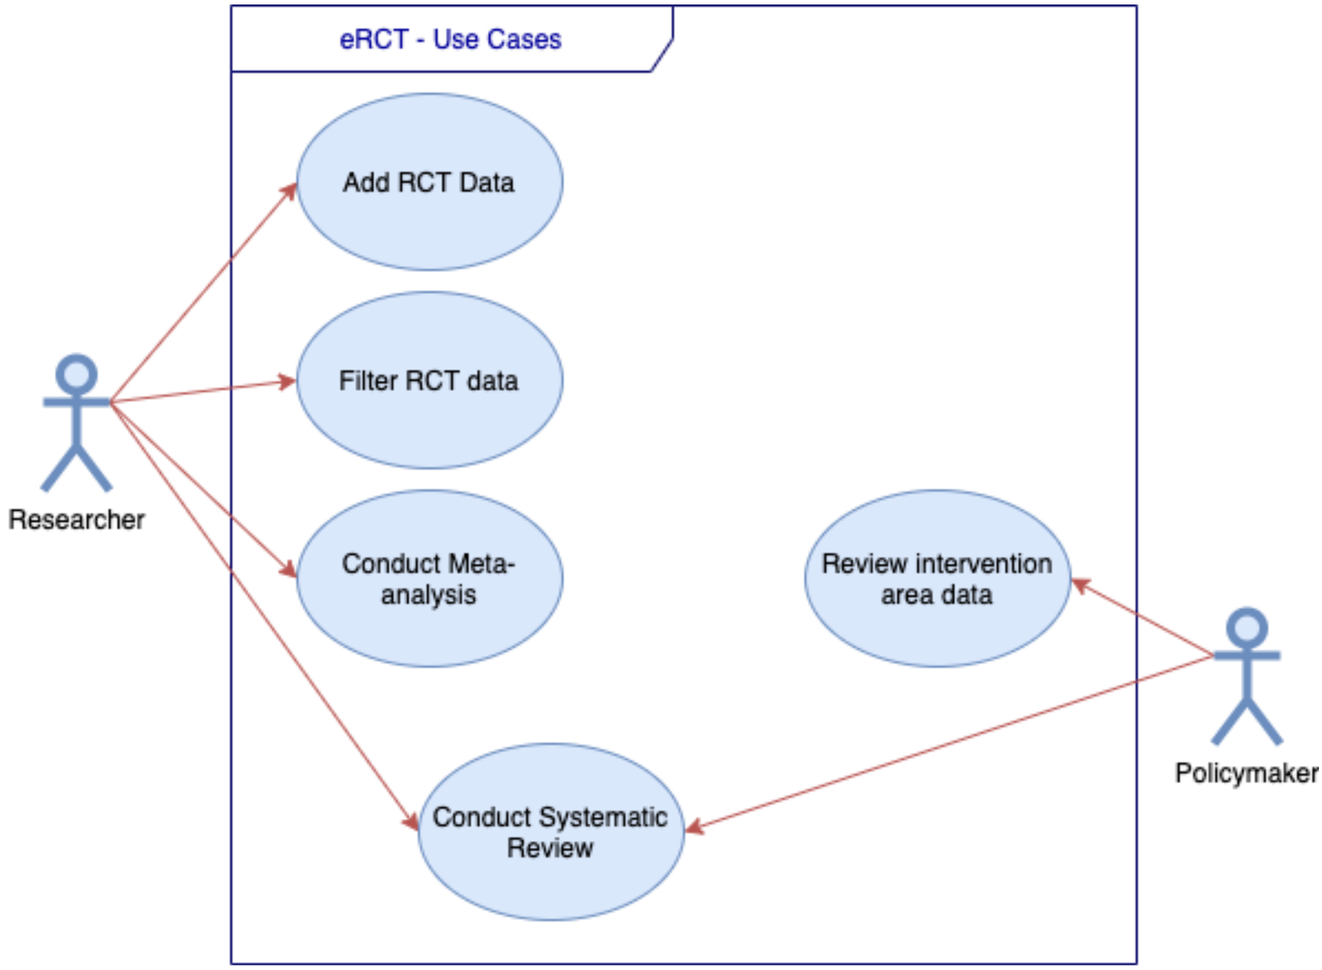


*Figure 4: Use case diagram for the ontology*

While use cases in ontology engineering often tend to be more open ended than traditional databases of software design (Gangemi and Presutti 2009), a number of generic use cases can be outlined here.

**Add RCT Data**

For researchers, a key element of the ontology is that it is open and replicable. Any study that is being conducted in the fields of social science or international development should be able to be expressed in the structure of the ontology, and it should be the case that data structured in this way can be contributed to linked knowledge graphs that are accessible by the research community.

**Filter RCT Data**

Researchers should be able to filter relevant RCT data for review, or for background research for developing new RCTs. To fulfil this requirement, researchers should be able to filter by contextual variables that provide analytically useful results. Questions related to this use case might include retrieving those RCTs that take place in those countries with the highest rates of teenage pregnancy, or those that are concerned with microfinance in those countries with the lowest levels of bank account ownership.

**Systematic Review**

A key element of conducting practical research is the systematic review: reviewing outcomes from related RCTs and drawing conclusions about the effectiveness of interventions in a given context. Researchers should be able to retrieve these outcomes from the ontology for review.

**Conduct Meta-Analysis**

In some cases, researchers need to go further than systematic review and pool treatment effects to provide a more quantitative outlook of the performance of a given intervention. This use-case defines that users should be able to access the quantitative output of RCTs in this area to enable this.

**Review Intervention Area Data**

The data produced in this area is similarly of interest to policymakers. Policymakers should be able to retrieve RCT data that can indicate which intervention policy to pursue. A typical question for this use case might be what the most effective intervention for increasing uptake in child vaccinations is, or what is the most effective intervention for increasing school attendance in disadvantaged areas.

## 3.4 Competency Questions

Having outlined the use case and domain for this project, it is advisable to compose a series of competency questions. Derived from the tradition of enterprise modelling, competency questions define a set of conditions that can be answerable once the ontology has been completed (Karima et al. 2017) and characterise the problems that the ontology is able to solve (Grüninger and Fox 1995). The intuition for this comes from the fact that ontologies are often considered models for a domain, but with use cases that may not be as tightly specified as traditional databases or software design – thus, this work requires a "Generic Use Case", a generalisation of potential use cases that serve as examples for domain modelling (Gangemi and Presutti 2009). Ideally, capturing domain tasks as competency questions involves producing a set of questions that an expert might wish to learn from a knowledge base: these questions thus serve as a guide to mould the ontology towards the needs of the users.

The framework used to define and evaluate the competency questions below based on the requirements of the ontology users is explained in detail as part of the evaluation framework in chapter 5.

**CQ1.** What RCTs in the East African Community (EAC) were stratified by gender?

**CQ2.** What RCTs in financial inclusion took place in countries with national savings rates of below 5%

**CQ3.** What RCTs in education took place in countries considered "highly developed" by the human development index (HDI)?

**CQ4.** What were the outcomes of cash transfer interventions in those countries considered "least developed" by the human development index (HDI)?

**CQ5.** What are the treatment effects for home tutoring interventions on standardised math scores for those countries with low levels of educational achievement?

**CQ6.** What were the outcomes of public health interventions in countries in the bottom 10% doctors per capita?

**CQ7.** What was the pooled treatment effect of microloan interventions on investment for those countries in the bottom 20% of GDP per capita?

**CQ8.** What proportion of studies in agricultural productivity could be considered successful?

## 3.5 Identification of Modules

Consistent with the modular approach outlined in this section, one of the primary tasks of the approach outlined in this paper is to identify modules that are consistent with user requirements and the existing structures of the data. Due to the size of the ontology, the dissertation does not outline every module of the ontology here, but rather discuss and focus on the key modelling choices taken in the design of the ontology. A full diagram of the ontology can be found in appendix A.3, describing the relationship between each of the modules outlined in the ontology. Furthermore, the full OWL expression of the ontology can be found in appendix A.6

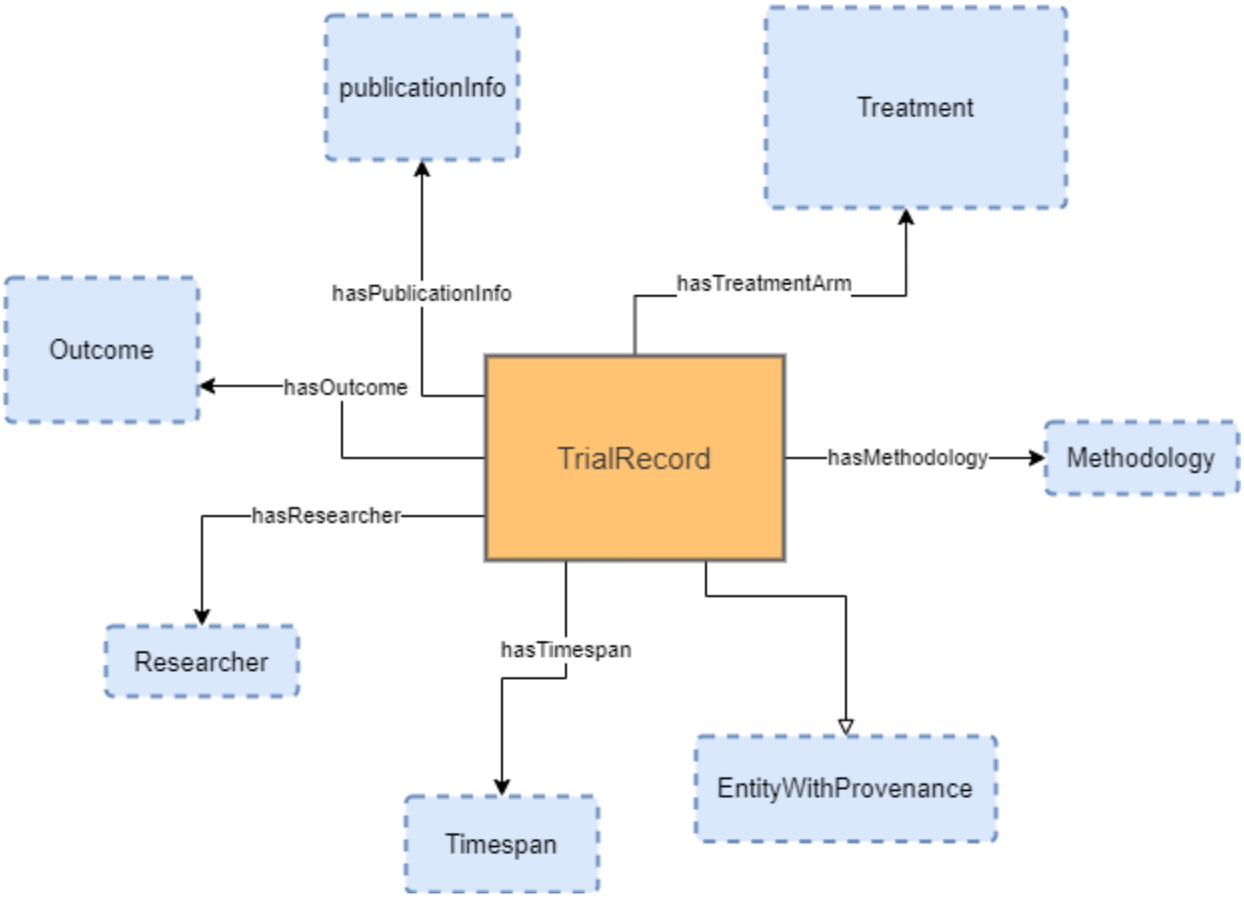


*Figure 5: Diagram for the TrialRecord module.*

*Consistent with the practice outlined in Shimizu et al. (2020), orange boxes represent classes and dashed blue boxes are classes which are part of another module. Black arrows represent the predicate relationship between the subject and object, while white arrows indicate a subclass relationship.*

The central aspect of the ontology model is the trial record, represented in figure 5 above. This key module is connected to nearly all other modules of the ontology, and the trial record object itself is a subclass of EntityWithProvenance, which includes all provenance information for the given RCT. This design follows that of Shimizu et al. (2020) and Mavergames et al. (2013). In the former, the ontology is built around an AgentRecord model that is further described by modules that capture the temporal extent of the agent, historical events that the agent took place in and similarly derives from an EntityWithProvenance module that describes the provenance information of the agent. In the latter, the central PICO module is more simply described by population, outcome and intervention modules, each of which describe the corresponding aspects of a medical RCT. Similarly, this ontology incorporates outcome and treatment modules, albeit describing different data and variables.

The schema above does not describe the axioms between modules. Figure 6 below includes a list of OWL axioms that describe the constraints between TrialRecord and other modules included in the ontology.

$$(1)\ TrialRecord \ \sqsubseteq EntityWithProvenance$$
$$(2)\ TrialRecord \ \sqsubseteq = 1 hasTimespan.Timespan$$
$$(3)\ TrialRecord \ \sqsubseteq \geq 0 hasOutcome.Outcome$$
$$(4)\ TrialRecord \ \sqsubseteq \geq 1 hasResearcher.Researcher$$
$$(5)\ TrialRecord \ \sqsubseteq \geq 0 hasPublicationInfo.PublicationInfo$$
$$(6)\ TrialRecord \ \sqsubseteq = 1 hasTreatment.Treatment$$
$$(7)\ TrialRecord \ \sqsubseteq = 1 hasMethodology.methodology$$

*Figure 6: OWL Axioms associated with the TrialRecord module*

These axioms are expressed in OWL and were encoded using Protégé (Musen 2015), a platform commonly used for building and maintaining OWL ontologies. Axiom (1) describes the subclass relationship between EntityWithProvenance and TrialRecord. (2) (6) and (7) describe existential relationships where TrialRecord contains one of each Timespan, Treatment and Methodology. (4) allows for the TrialRecord to have many researchers while requiring at least one, reflecting the fact that most RCTs are conducted by research teams. Axiom (5) allows for zero or more publicationInfo modules, reflecting the fact that the results of some RCTs tend to be published in multiple journals or other formats, and in some cases, may not be published at all, while (3) similarly allows for multiple outcomes and no outcome in the case that the RCT has yet to be completed.

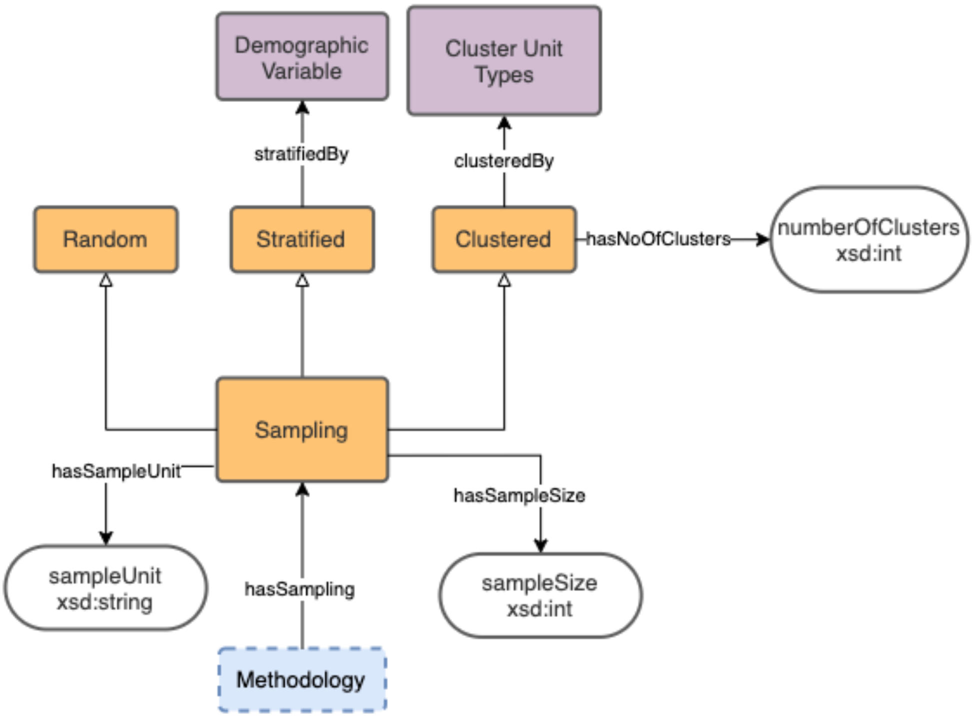


*Figure 7: Diagram for the Sampling module.*

*Representation as in figure 5, with purple boxes indicating that the object forms part of a controlled vocabulary, while white oval boxes represent literal values.*

The sampling module, referred to in figure 7 above, is a novel pattern contribution made as part of this research. It primarily focuses on the three main types of sampling found in RCTs: simple randomised sampling, randomised sampling through stratification across a given set of demographic variables and cluster sampling, whereby participants are divided into clusters based on accessibility or a similar criterion. Each of these are subclasses of the sampling object, in which the total sample size is defined in addition to the sample unit, which in most cases will be individuals but, in some cases, might be schools or villages depending on the type of intervention. In addition to this, clustered-type sampling objects have the additional property of defining the number of clusters in included in the study.

The demographic variable and cluster unit type objects are examples of a *controlled vocabulary,* a concept derived from Shimizu et al. (2020). These are objects that can take a standardised set of values that are defined in metadata rather than being structural elements of the ontology itself. This is a conscious design decision made in the areas of the ontology that are most likely to require expansion in the future: by not using subclassing in these areas, the types included in the controlled vocabulary can expand, increasing the query options available to users while not requiring an update or redesign of the entire ontology itself. This decision is made in response to warnings to “beware of the hierarchy”, an observation by Pernisch et al. (2021) of the wide ranging impact updates to an ontology can have on downstream semantic applications that they rely on, with particular concern levied towards those ontologies that use a high degree of subclassing.

In this instance, the random, stratified, and clustered types of sampling are unlikely to change in the future, being a fundamental feature of any randomised study. The demographic variables that an RCT might be sampled on can change, however. Furthermore, interviews conducted as part competency question design indicated it would be difficult the determine the entire universe of possible demographic variables that might be used for this during the design phase of the ontology. It was similarly difficult to determine the diversity of cluster unit types that users might encounter while using the ontology. Therefore, the choice was made to include each of these as controlled vocabulary types.

Figure 8 below shows a snippet of the OWL ontology encoded as RDF, produced through protégé. The full OWL expression of the ontology can be found in appendix A.6, in addition to this project’s Github repository.

```
 <!-- http://www.semanticweb.org/ERCT#Sampling -->

    <owl:Class rdf:about="http://www.semanticweb.org/ERCT#Sampling">
        <rdfs:subClassOf>
            <owl:Restriction>
                <owl:onProperty
rdf:resource="http://www.semanticweb.org/ERCT#hasSampleSize"/>
                <owl:qualifiedCardinality
rdf:datatype="http://www.w3.org/2001/XMLSchema#nonNegativeInteger">1</owl:qual
ifiedCardinality>
                <owl:onDataRange
rdf:resource="http://www.w3.org/2001/XMLSchema#int"/>
            </owl:Restriction>
        </rdfs:subClassOf>
        <rdfs:subClassOf>
            <owl:Restriction>
                <owl:onProperty
rdf:resource="http://www.semanticweb.org/ERCT#hasUnit"/>
                <owl:qualifiedCardinality
rdf:datatype="http://www.w3.org/2001/XMLSchema#nonNegativeInteger">1</owl:qual
ifiedCardinality>
                <owl:onDataRange
rdf:resource="http://www.w3.org/2001/XMLSchema#int"/>
            </owl:Restriction>
        </rdfs:subClassOf>
    </owl:Class>
```

*Figure 8: Snippet of Sampling Object within the OWL Ontology*

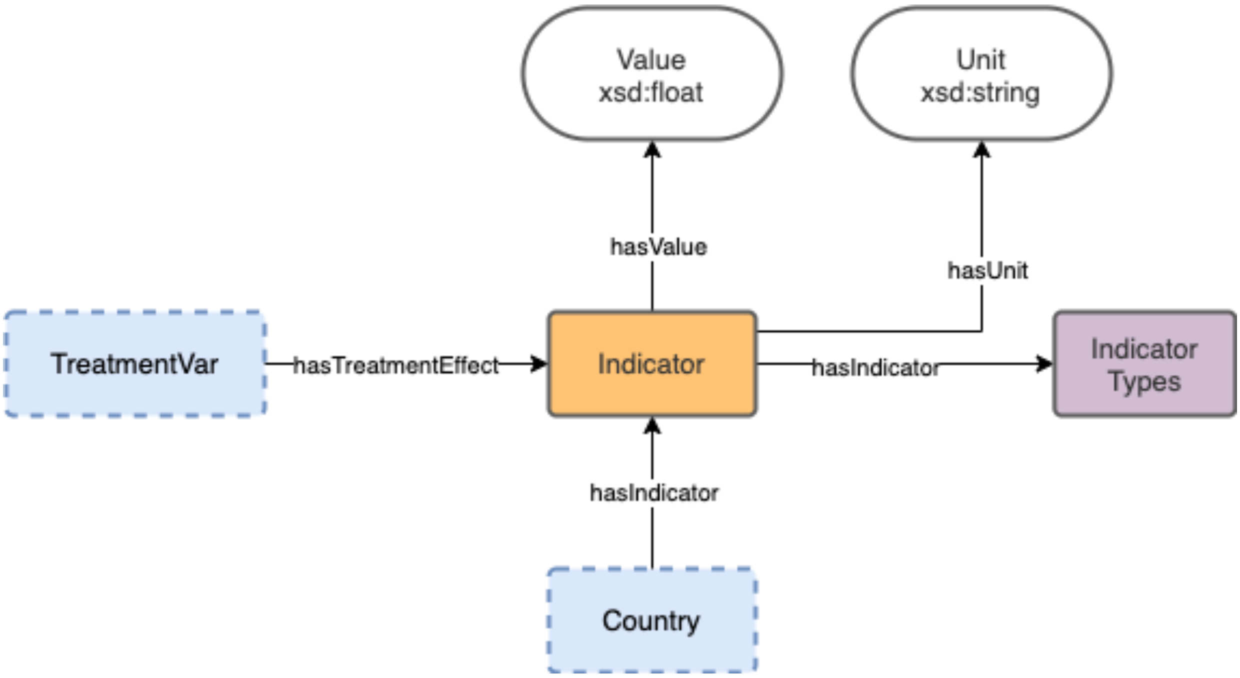


*Figure 9: Diagram for the Indicator module*

*Representation as in figure 7.*

The indicator module represented in figure 9 above is another novel addition of this research. This module focuses on delivering quantitative information that can arise from a diversity of sources. In this ontology, the indicator module is used to represent the quantitative results of an RCT contained within a treatment effect, in addition to indicator values that describe the development progress of a country or any other defined entity. The value and unit objects describe the numerical value and unit of measurement of the indicator respectively.

The indicator types object is another representation of a controlled vocabulary. As in the case of the sampling module, interviews with domain users revealed that the range of types of results that could arise as result of RCTs would be very difficult to ascertain during the design phase. Additionally, it was identified that the types of results that RCTs provide would evolve over time with the discipline as researchers move towards new topics and explore different avenues of research. Thus, a controlled vocabulary was implemented in order to allow for this area of the ontology to more easily expand.

## 3.6 Creation of OWL Files in Protégé

Following the design of the ontology, the outlined structure was created in Protégé and expressed as OWL files. These files represent the described axioms in machine readable form. These are available in the ERCT github repository, a link to which is available in appendix A.4.

# Chapter 4

# Implementation

Having discussed the overall design of the ontology, this chapter now considers the implementation of the ontology, including data uplift and access via a SPARQL endpoint. The purpose of this exercise is to provide a platform through which the ontology can be evaluated, allowing the ontology to be assessed against its intended purposes.

## 4.1 Data Pipeline

The aim of this ontology is to allow users to access relevant outcomes for systematic review to assess outcomes across methodological directions, policy areas or research areas. As such, consistent with previous approaches to such resources, this approach calls for a *demonstrator hub,* which allows for the ontology to be demonstrated at least initially with data from a limited number of databases and allowing it to grow over time as collaborators and contributors begin integrating with the project (Shimizu et al. 2020). This data hub provides a vehicle for evaluation of the ontological design approach, which can be checked against the competency questions outlined in chapter 3 that informed its design.

Figure 10 below demonstrates the proposed data pipeline for the initial demonstration hub of the ERCT ontology. This involves processing data from a variety of sources, collating this into an amenable format with a python based pre-processing step, uplifting this data from a tabular to graph format with the R2RML engine and finally storing this in a query engine accessible through SPARQL queries or similar.

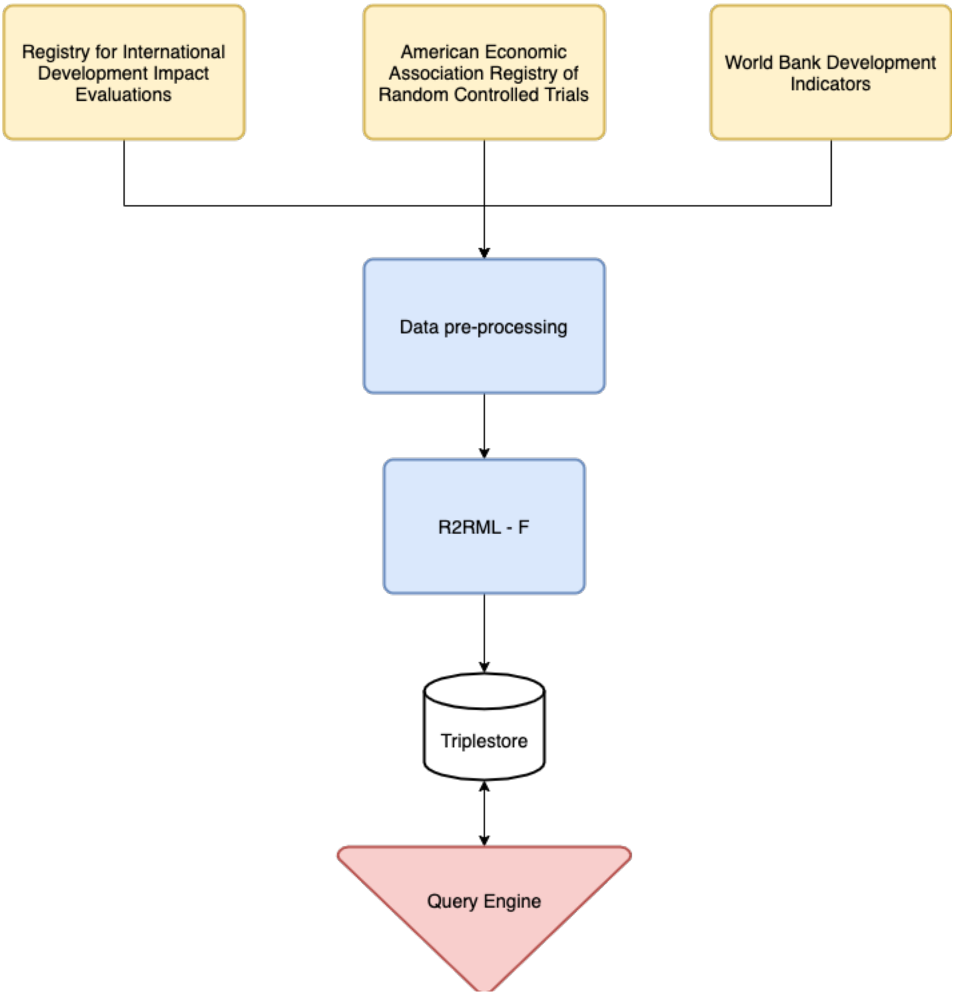


*Figure 10: The Data Pipeline Process for the Ontology*

## 4.2 Input Data

There are two main data types of data sources for this project: trial data, which includes variables such as the methodology and outcomes of a given randomised control trial, and contextual data, which provides key macro-level economic and societal indicators for the given geography that the trial took place in. The trial data was chosen to provide the necessary variables to conduct a systematic review or metanalysis. The contextual data was included so that researchers could conduct a filtering step, allowing them to choose the RCTs that appear in each context. Additionally, including contextual data serves as a starting point for further ontology development in the area of economic, social and international development terms and indicators, an area of ontology development that has, to the knowledge of the author, been sparsely researched.

For the former, data is sourced from the from the AEA RCT registry (American Economic Association 2021) and the RIDIE registry (International Initiative for Impact Evaluation 2021). The most up-to-date registry from the AEA is provided via a csv download. The RIDIE database, data is only accessible via their search function, with the data displayed as html through a website. A web scraper was therefore produced to convert this data into a tabular format[2]. For the contextual data, the data is sourced from the World Bank development indicators, accessible via a stable and consistent API that is widely used by data providers and development and research organisations around the world (World Bank 2021).

| | Size | Key Variables | Lowest Granularity |
|---|---|---|---|
| AEA RCT Registry | 4707 rows, each RCT as row | Outcomes, Sampling, Research area etc. | RCT Level |
| RIDIE Registry for International Development Impact Evaluations | 5980 rows, each RCT as row | Outcomes, Sampling, Research area etc. | RCT Level |
| World Bank Development Indicators | Indicator dependent, at most 211 | Gross Domestic Production, Human Development Index, Proportion of Population with access to safe drinking water | Country Level |
| ERCT | 378,243 triples | As above | RCT Level |

*Table 1: Summary of Data Included in Uplift*

[2] The was produced in Python, and the code is available in Appendix A.4

## 4.3 Data Processing

The primary challenge with the data processing step for this pipeline is developing a reproducible and sustainable method for the discussed tabular data to be reconciled into an RDF format. There are several methods that have been proposed to achieve this. The simplest of these is the RDB2RDF mapping language (Arenas et al. 2012), which simply creates an RDF node based on each record with the column representing the propertyType and the table cell representing the corresponding value (Arenas et al. 2012). While simple to implement, the process is not nuanced enough to express the trial and contextual data: it does not allow for control over which rows are uplifted in a situation where not all rows are relevant, does not allow us to differentiate predicates for specific rows or perform any data cleaning or processing on selected values. A partial solution to these issues is R2RML, which allows for a declarative approach where specific rows can be selected via SQL and then assigned predicates either specified in mapping or via a value in the logical table being processed (Das et al. 2012). The uplift method is completed by R2RML-F, an extension to R2RML, which allows for some manipulation of field data through the addition of ECMAScript based functions that can be integrated into the R2RML mapping (Debruyne et al. 2016).

Even with these extensions, there are some elements of the tabular data that need to be pre-processed to be used by R2RML. Columns that contain comma separated data need to be separated out into individual columns so that they can be processed into separate object values by R2RML, and some data standardisation must be completed, such as converting country and place identifiers in the trial data to the ISO-3366 and ISO-3366-2 standards respectively. This allows for the data to be linked with the World Bank development indicator data.

Figure 11 below shows a snippet of the R2RML implementation used to map variables from the various datasets into the RDF format, expressed in triples. This snippet identifies sample size within the tabular data and associates with the appropriate object in the ontology. The full R2RML mapping is available in the project’s Github repository, which can be found in appendix A.4.

```
<#SamplingMapSampleSize>
    a rr:TriplesMap ;
    rr:logicalTable [
        rr:sqlQuery """SELECT RCT_ID, OBS FROM merged"""
    ] ;
    rr:subjectMap [
        rr:template "http://www.semanticweb.org/Sampling/{RCT_ID}" ;
    ] ;
    rr:predicateObjectMap [
        rr:predicate erct:hasSampleSize ;
        rr:objectMap    [
        rrf:functionCall [
                rrf:function <#ExtractDigits> ;
                rrf:parameterBindings (
                    [ rr:column "OBS" ]
                ) ;
            ] ;
        ]
    ] ;
    .
```

*Figure 11: Snippet of the uplift mapping used to uplift the data to RDF, implementing R2RML-F*

## 4.4 Triplestore and Querying

Following the production of the new RDF dataset according to the ontology, the data must be uploaded to a triple store. For this purpose, GraphDB is implemented, a graph database developed by Ontotext (2021) that allows for localised querying via SPARQL (Prud'Hommeaux and Seaborne 2008), a query language that allows users to query data stored as RDF triple stores

# Chapter 5

# Evaluation

Following the design and implementation of the ontology, this chapter now considers the evaluation of the ontology. In particular, this chapter focuses on the nature of the proposed evaluation and the steps taken to ensure that the evaluation is a meaningful reflection on how the ontology reflects the competency questions and use cases outlined in chapter 3. The proposed methodology for this study was approved by the SCSS ethics committee[3].

## 5.1 Structure and Design of Evaluation

The testing and evaluation of ontologies is often a significant challenge, as the purpose of most ontologies is to allow for the posing of novel questions, which by definition cannot have existing gold standard answers to which to compare. This is further complicated by the domain of RCTs in the social sciences – there exists limited material at present that provides a well-established set of principles to guide researchers in the systematic review process. This creates a deficit of previous work from which to build a framework to assess the ontology might have been derived. While one of the most definitive guides in the field cites the necessity to increase the number of systematic reviews and meta-analysis of existing work, it does not describe a framework or set of principles through which this might be achieved (Banerjee and Duflo 2017). Furthermore, a user evaluation of this platform would be difficult to achieve, given that the end-user of this assessment is intended to be those working in policymaking, but with the ontology used as an intermediary tool for knowledge access.

[3] The ethical approval can be viewed in appendix A.2

Instead, the evaluation of the competency questions is developed as part of the design phase. In order for these questions to be as reflective of the needs of the use case as possible, two methods are employed. First, interviews were conducted with three experts with a wide breadth of experience in conducting, reviewing and managing RCTs in a variety of areas in the social sciences, including economics and international development. This data allowed for the shaping of competency questions from input directly elicited by experts for whom the ontology is intended. Secondly, the processes and practices that have been implemented in existing systematic reviews and meta-analyses are reviewed, with particular attention given to the selection process and output evaluation of the studies reviewed. While a standardised codified procedure for such reviews has yet to emerge, this step will allow us to create competency questions that reflect the most common practices adopted in existing reviews

## 5.2 The Interview Process

The interview section of this process was designed to collect data from of those experienced in the practice of RCTs on the key areas of interest of their area of study. The interviews were designed to be semi-structured, with a set of guiding questions[4] used to allow the interviewee to express their thoughts on the most pressing empirical questions in their area of research, while also allowing for flexibility to pursue a particular line of questioning if it showed promise in terms of relevant data. In this way, they are designed as *focused interviews*, where the scope of topics is maximised to allow the expression of points of view that have not been anticipated (Hopf 2004).

Participants were sampled through simple random sampling, drawn from all researchers who have submitted their RCT on the AER registry. The reasoning behind this is that those who submitted their research to this forum tend to be the most regular practitioners of RCTs as a method of research, and therefore appropriate interviewees for this research. This sampling strategy and a sample size of five was decided based on an equilibrium of research needs and sampling concerns (Robinson 2014). Unfortunately, the response rate was low: of the 250 researchers contacted, only three agreed to interview, while the aim of the research was to

[4] The full list of questions used is available in appendix A.1

interview five. Given the difficulty in finding researchers willing to interview, future work in this area might focus on increasing the number of interviewees through alternative networks.

| Interviewee 1 | Interviewee 2 | Interviewee 3 |
|---|---|---|
| Education economist working in a major international development organisation. Has overseen and managed RCTs across several developing countries and published journal articles based on results. | Economist in a European government agency responsible for using RCTs as part of validating and developing policy. Leads the implementation of RCTs and develops policy recommendations from results. | Economist and legal scholar based in North America. Uses RCTs as part of generating information on the effects of the legal system in the country of interest on a variety of variables. |

*Table 2: Profile of interviewees*

All interviews were recorded and transcribed in line with the ethical approval received. The material from these interviews was then used to shape the competency questions, with particular attention given to the features of RCTs that the researchers look to when evaluating their quality and impact.

From the interview questions outlined in appendix A.1, the most insightful questions were those focused on methodology, namely questions one and four and their subcomponents. These questions elicited the most pertinent information on how the researchers critically engage with RCTs, and which methodological variables are sought when reviewing evidence in the area. Questions 1.a and 1.b provided a clear focus on sampling as one of the key areas that researchers focus on in review. Similarly, question 4b provided an opportunity to access a variety of papers focused on the design and implementation of RCTs, such as Bannerjee and Duflo's *Handbook of Field Experiments* (2017), which provided a detailed overview of the lifecycle of RCTs in economics and the social sciences. Questions 2 and 3 were informative but outside of question 2.a, provided less data directly relatable to the design of the ontology. From the experience of this dissertation, future iterations of this work would focus on rephrasing these questions to focus more on potential types of treatment results and effects, in addition to understanding categories and subcategories of research interest that could be incorporated into the ontology.

## 5.3 Design Decisions Derived from Interviews

Several design decisions arose from these interviews. The primary outcome was insight into where in the ontology to adopt types and sub-classes. As discussed in section three, this is a significant design choice for any ontology; the former allows for more flexibility but can sometimes come at the cost of intuition, while the latter can create more intuitive connections between objects in the ontology while requiring a more extensive redesign in the event that the ontology needs to expand. From all three interviews it became clear that the total universe of possible intervention treatment types, demographic variables used for stratification, cluster unit types for cluster sampling, result types for treatment effects and experimental design types would be too extensive to allow for sub-classing. Instead, as with Shimizu et al. (2020), a controlled vocabulary was adopted for these objects, with a guide included with the ontology indicating the types that can be used and where.

A second design decision influenced by the interviews was the structure of the sampling module. As part of the interviews, the interviewees were asked for the elements of RCTs they look for first in academic papers. A common response for this was to understand how the RCT was sampled, with particular attention given to how the authors designed the stratification of the sampling. The ontology therefore reflects this by including a substantive module on sampling, including the possible demographic variables that the sample stratified on and the unit and total number of units within this sample.

## 5.4 Review of Existing Systematic Reviews

For this process, the steps outlined in the *search strategy* phase of Waddington et al. (2012) was followed as much as possible; more specifically, searches were conducted on a variety of standard databases, with particular attention to multi-disciplinary databases. A screening and eligibility process was conducted on the papers found with this method, filtering out those that are not relevant to RCTs within the social science domain. The process departs from Waddington et al. in two important ways. First, concerns about publication bias were not as high a concern given that this process was concerned with screening systematic reviews rather than trials themselves. Second, the scope and resources available to this study is smaller than those that this guidance is aimed towards. Therefore, the number of studies chosen for review had to be limited to a sample of those eligible following screening.

Following the identification and filtering phases, a selection of five existing systematic reviews in empirical economics were chosen for review. Those chosen for review were Fryer (2017), Baird (2013), Edwards (2013), Bastagli et al. (2016) and Saavedra and Garcia (2012). These papers were chosen for a variety of reasons; they offer a mixture of both systematic review and meta-analysis approaches, cover a variety of different topic areas, and come from policy and academic focused sources. These five review papers cover the areas of health economics, microfinance and cash transfers and education economics. Each paper was compared for common methodological processes.

A process common to each paper was a **filtering step**. This involved selecting a set of papers to be included in the systematic review based on some criteria that they fulfil, such as those employing a particular type of empirical methodology or those that were conducted in a particular area of research or policy interest. For example, in the case of Fryer (2017), almost one thousand studies were initially chosen as potential candidates based on the methodological condition that they use a "verifiably random procedure to assign participants to treatment and control groups in a non-laboratory environment" (Fryer 2017:5). This number was further reduced to 199 by limiting the studies to those conducted in "highly developed" countries with standardised reading or math outcomes, a classification measured through the human development index, an amalgamation of that country's life expectancy, mean and expected years of school and gross national income per capita. Finally, the authors divide the sample of studies into interventions concerning early childhood, school-based interventions, and home-based interventions for analysis. Similarly, Bastagli et al. (2016) filter their systematic review based on the impacts of cash transfers in six relevant areas of policy, such as education, investment, or savings.

Secondly, authors tend to conduct a **data collection step**. Studies diverge slightly in this area, with some simply gathering the outcomes of each study considered and commentating on the quality and extent of their results for a given policy or research area (Bastagli et al. 2016). Other studies perform a more intensive data collection step with the intention of meta-analysis; both Baird (2013) and Fryer (2017) collect treatment effects directly from results tables, with Fryer collecting the impact on standardised math and reading scores for students for a range of over 40 explanatory variables[5].

Finally, authors implement **analysis step**. For those studies that are more systematic review oriented, this involves an intensive discussion of the outcomes of the studies chosen in the previous steps, with particular attention to their implications for effective interventionist policy (Bastagli et a. 2016, Edwards 2013). For those studies that are more oriented towards data focused meta-analysis, the relevant treatment data compiled from tables are collated into a model which can inform provide a more accurate estimate of the treatment effect for a particular intervention (Baird 2013, Fryer 2017).

## 5.5 Design Decisions Deriving from the Review of Systematic Reviews

Apart from contributing to the design of the evaluation, the review of reviews led to a number of design decisions for the ontology. One of the key decisions that derived from this analysis was in response to quantitative meta-analysis (Baird 2013, Fryer 2017), which seeks to pool treatment effects to discover an aggregate effect of an intervention on the population. In response to this, a treatment effect module was included in the ontology, which includes a description of the value and unit type of the treatment, in addition to a limited vocabulary object indicating the type of result that was produced. Ultimately, it is hoped that this module

[5] These impacts were standardised as the change in standard deviations of test score outcomes based on a given intervention.

can be used to pool interventions of similar type in order to improve the process of meta-analysis. However, as outlined later in this chapter, there are issues surrounding the sourcing of data for this module. Possible future work that could alleviate this is included in chapter 6.

The second design decision that was influenced by this review was the inclusion of standardised country codes and contextual data. Influenced by the filtering step common to all of the papers included in the review, linking the data to key economic and social development indicators provided by the World Bank allows for those conducting systematic review to quickly filter to those results relevant to a particular context when retrieving RCT data. Some systematic reviews for example, might be relevant to highly developed economies, while others might focus on those countries with low levels of some indicator, such as bank account ownership or level of education.

## 5.6 Tying Competency Questions to Evaluation

Following the expert input and review of existing processes for systematic review, the competency questions are divided into groups that cover each part of the lifecycle of the systematic review process. The intention of these competency questions is to accurately reflect the systematic review process, allowing an evaluation of which parts of the process the ontology can successfully deliver pertinent results, and if there are any deficits in a particular part of the review lifecycle that the ontology might improve on in the future.

Additionally, the evaluation criteria of each competency question are broken down into four distinct parts:

| | |
|---|---|
| **SPARQL** | This evaluation criterion indicates whether this question can be successfully implemented is due to the capabilities of the SPARQL query language. This is particularly pertinent for questions of a strong analytical nature – if a question calls for a high degree of analytical processing of output data to be answered, it is more likely that SPARQL |

will not be suitable and that a downstream language like python would need to be used in conjunction with it.

**Data – RCT** This evaluation criterion indicates if the question can be answered based on the available RCT data. Due to the limitations of data collected by the AEA Registry and RIDIE, there may be cases where some data is not available to satisfy the structure of the ontology as it was informed by competency questions. Questions that require detailed quantitative RCT data are more likely to fail this criterion.

**Data – Contextual** This criterion indicates if the question can be answered based on the availability of the contextual data, which is development indicator data derived from the World Bank. Questions with distinctive filtering requirements are more likely to fail this criterion, as a country level development indicator may not exist to provide the required information.

**Ontology** This criterion indicates if the question can be answered based on the structure of the ontology. If the ontology does not contain the necessary relations or objects to satisfy the question, it fails this criterion.

## 5.7 Evaluation of Competency Framework

To evaluate the ontology based on the competency questions set out in chapter 3, a set of SPARQL queries were created to represent each of the questions posed. While this is not an exhaustive analysis of the potential of the ontology, it represents a good demonstration of how the ontology performs with a given use case. Each of the competency questions were created within the realm of the steps identified as commonly part of systematic review from sections 5.4 and 5.5, with each competency question representing a common query within that step. Each of the competency questions are evaluated against the criteria identified in section 5.6 above. All of the queries created to represent these competency questions can be viewed in appendix A.5.

The results can be seen from tables 3 to 6. A clear trend can be seen in how the ontology operates: while the ontology is clearly effective in filtering specific RCTs for analysis, it performs less well in providing specific treatment results regarding interventions, or in amalgamating those treatments that might be considered successful or unsuccessful. As can be seen in the breakdown of the evaluation criteria, this issue largely arises from a lack of complete RCT data. This section will proceed by describing the outcomes of each of competency question area across all evaluation criteria.

### 5.7.1 Filtering Step Questions (CQ1-CQ3)

The ontology performs well for each question associated with the filtering step of the systematic review process, providing relevant results based on macroeconomic and methodological variables. Within the context of CQ1 – CQ3, the three competency questions associated with this filtering step can be fulfilled by all key functionalities associated with the ontology. The SPARQL language provides the required functionality to retrieve and filter results by all relevant variables, allowing for results to be filtered both in a binary fashion as in gender in CQ1, or in a more ranked manner is in CQ2. The required data from the RCT databases as well as the contextual data from the world bank both provide the necessary data required to answer these queries. Finally, the ontology itself provides the necessary logical axioms to allow the user to specify the results required.

| ID | Competency Question | SPARQL | Data | | Ontology |
|---|---|---|---|---|---|
| | | | RCT | Contextual | |
| CQ 1 | What RCTs in the East African Community (EAC) were stratified by gender? | Y | Y | Y | Y |
| CQ 2 | What RCTs in financial inclusion took place in countries with national savings rates of below 5%. | Y | Y | Y | Y |
| CQ 3 | What RCTs in education took place in countries considered "highly developed" by the human development index (HDI)? | Y | Y | Y | Y |

*Table 3: Competency questions for the filtering step*

*Focused on the filtering stage of systematic review, each column represents a functionality associated with the ontology, while each row on these columns indicates whether that functionality can satisfy the competency question posed.*

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>


PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
select  ?country ?countryCode ?DemVar ?countryName where {
   ?country erct:hasISOAlpha2 ?countryCode.
   ?country erct:hasName ?countryName.
   FILTER (?countryCode IN ("RW", "KE", "UG", "TZ", "BU", "SS", "SD") ).
   ?Stratified a erct:Stratified;
       erct:stratifiedBy ?DemVar.
       FILTER regex(?DemVar, "gender", "i").
}
```

*Figure 12: SPARQL Query Representing Competency Question CQ1*

| | country | countryCode | DemVar | countryName |
|---|---|---|---|---|
| 1 | http://www.semanticv STUDY-ID-5356d701c1052 | "RW" | "gender" | "Rwanda" |
| 2 | http://www.semanticv 0006837 | "RW" | "gender" | "Rwanda" |
| 3 | http://www.semanticv 0000161 | "RW" | "gender" | "Rwanda" |
| 4 | http://www.semanticv 0002565 | "RW" | "gender" | "Rwanda" |
| 5 | http://www.semanticv 0002214 | "RW" | "gender" | "Rwanda" |
| 6 | http://www.semanticv STUDY-ID-5c9848936449a | "RW" | "gender" | "Rwanda" |
| 7 | http://www.semanticv 0001030 | "RW" | "gender" | "Rwanda" |
| 8 | http://www.semanticv 0000352 | "RW" | "gender" | "Rwanda" |

*Figure 13: Output from CQ1, showing RCTs from the East African Community States stratified by gender*

### 5.7.2 Data Collection Step Questions (CQ4 – CQ6)

In the data collection process, the implementation is more mixed. The ontology performs when providing the summarised outcome of studies that researchers have included as part of their submission to the original RCT database that data was sourced from. This output is beneficial to systematic reviews that compare outcomes of studies qualitatively, such as Bastagli (2016), as it allows for these reviews to categorise and discuss the results that have been provided in the database.

For those studies concerned more with meta-analysis that pools treatment effects such as Baird (2013) and Fryer (2017), the ontology performs poorly. As indicated in table 4, this is singularly a data availability issue: most registries do not categorise treatment effects in a way that is conducive to uplift, with many studies including them as blocks of text in short summaries of the research output, while others do not include them at all. Until data providers enforce standards as to how is data is collected and received by researchers, the only solution to this issue is to source the data from other sources using crawlers and classifiers. While outside the scope of this dissertation, possible methods of achieving this are discussed in chapter 6.

| ID | Competency Question | SPARQL | Data | | Ontology |
|---|---|---|---|---|---|
| | | | RCT | Contextual | |
| CQ 4 | What were the outcomes of firm productivity focused interventions in those countries considered “least developed” by the human development index (HDI)? | Y | Y | Y | Y |
| CQ 5 | What are the treatment effects for home tutoring interventions on standardised math scores for those countries with low levels of educational achievement? | Y | N | Y | N |
| CQ 6 | What were the outcomes of public health interventions in countries in the bottom 10% doctors per capita? | Y | Y | Y | Y |

*Table 4: Competency questions for the data collection step*

*Competency questions as part of a framework for evaluating the lifecycle of the systematic review process, focused on the data collection stage of systematic review. Columns and rows as in table 3.*

### 5.7.3 Review Questions (CQ7 – CQ8)

The implementation performs poorest on the review phase, with neither of the competency questions posed answerable by the ontology. As with CQ4-CQ6, this is primarily a data problem: without some way to capture and classify quantifiable treatment effects that are included in the ontology, it remains impossible to pool treatment effects for analysis as in CQ7, on only the qualitative or summarised output of CQ8 can be assessed. Furthermore, the SPARQL query language encounters difficulty with CQ7. While the language is adept at fetching appropriate results and performing some basic aggregation tasks, the more complex task of pooling treatment effects is outside of the scope of the language and would have to be performed with a python script or similar. This would not necessarily limit the effectiveness of the ontology in this area should the treatment effect data become available, but would prevent any semantic application built on the ontology from exclusively relying on the query language to produce results.

| **ID** | **Competency Question** | **SPARQL** | **Data** | | **Ontology** |
|---|---|---|---|---|---|
| | | | **RCT** | **Contextual** | |
| CQ 7 | What was the pooled treatment effect of microloan interventions on investment for those countries in the bottom 20% of GDP per capita? | N | N | Y | N |
| CQ 8 | What proportion of studies in agricultural productivity could be considered successful? | Y | N | Y | N |

*Table 5: Competency questions as part of a framework for evaluating the analysis step*

*Competency questions as part of a framework for evaluating the lifecycle of the systematic review process, focused on the analysis stage of systematic review. Columns and rows as in table 3.*

# Chapter 6

# Conclusion

## 6.1 Summary of Findings

This dissertation developed and evaluated an OWL ontology to describe the methodology and outcomes of RCTs in the field of economics, the social sciences and international development. The ontology was developed and evaluated through the use of semi-structured interviews and a review of existing systematic reviews and meta-analyses.

From the evaluation of the ontology through the developed competency questions, it was found that the ontology satisfies approximately 70% of identified user need from the ontology. The key lesson from this evaluation was that while the ontology tends to perform the filtering step of systematic reviews particularly well, providing data on papers that are relevant to review, the ontology performs less well in providing and processing a wide range of quantitative outcome data. The primary reason for this is data: while the contextual data consistently allows for efficient filtering based on key development, economic and demographic statistics, the RCT data is lacking in quantitative outcome data. More specifically, key quantifiable outcomes used in meta-analysis such as regression output and treatment effects are not available. Furthermore, SPARQL has difficulty as a data processing platform in pooling quantitative data together – more advanced statistical analysis in meta-data-based reviews must be performed using downstream statistical software.

## 6.2 Primary Contributions

This makes contributions in several areas. First and foremost, to the best of the author's knowledge, this represents the first ontology that attempts to frame relevant data on the outcomes and methodology of RCTs in economics, the social sciences and international development in a way that is conducive to systematic review. As such, the evaluation in this paper serves as the lessons learned from this exercise and paves the way for such an ontology or similar ontologies to be developed further in the future. Given the nature of RDF and OWL, it also allows for future ontologies and data in the domain to link to this ontology, expanding the possibilities of semantic applications.

Secondly, the work contained in this paper represents the first attempt to uplift RCT data from multiple sources into a cohesive ontology, linking disparate RCT databases for the first time. This is furthered by the inclusion of World Bank development indicators within the ontology, linking contextual data to RCT data for the first time. Given that there are currently no ontologies used by the World Bank or any other key data provider in this area, this represents the first time that a framework has been provided that can provide one point of entry for each of these data pipelines. Through future work on harmonising these data sources, it is possible that a consistent SPARQL endpoint could be provided to allow semantic web applications to be built in this area.

## 6.3 Future Work

The key lesson learned from the evaluation in this paper is that the main factor preventing user needs from being fully fulfilled by this ontology is due to a lack of data. In particular, key quantitative data on treatment effects are not available on RCT registries – this does not allow for more quantitative focused meta-analysis to be completed on those studies included in these registries. By including this data, it could be possible to develop semantic solutions to systematic review that can perform more complex analysis, such as creating pooled treatment effects that can reveal "global" results of RCTs, providing a clearer picture as to whether perceived results are applicable across multiple settings.

One way to achieve this might be to extract the data directly from papers that are produced as a result of these studies. Given the volume of research papers and monitoring and evaluation studies produced by RCTs, this would not be an insignificant task. However, it is not without precedent – TableSeer (Liu et al. 2007), a tool developed as part of the Citeseer digital library (Williams 2014) can crawl and extract tabular data included in academic papers, while more recent solutions such as AxCell (Ruder et al. 2020) uses the Latex source code of the paper to classify and extract table results, with particular attention given to results from machine learning papers. The more difficult element of this task would be associating treatment results from these papers with the correct RCT entry from the ontology and presenting this data to the user in a functional manner.

Another issue observed in the evaluation is that SPARQL became unsuitable as an analytical language for more complex data tasks, such as pooling treatment effects. While this issue can be alleviated with a downstream analytics step using a more data focused language like Python, ideally future projects built upon the ontology would employ an entirely semantic web based solution. One possible way to alleviate this issue would be to use an extended version of SPARQL, such as Recursive SPARQL (Hogan et al. 2020). However, while this can be deployed on top of traditional SPARQL engines at present, basing the functionality of the knowledge graph off of an experimental extension would likely introduce future issues of compatibility and support.

One drawback to this approach is one common to almost all systematic review processes: extracting results only from those studies that have produced papers potentially introduces publication bias into any systematic review created from the data, which is where those papers that have significant positive results are more likely to be published by journals. A recent analysis of this issue in the area of empirical economics revealed that those studies that employ a difference in difference estimator in their methodology, commonly used in RCTs, tends to have a slightly higher incidence of p-hacking than other methods[6], though incidence of this problem

[6] In this particular study, it was found that only instrumental variable approaches has a higher incidence of p-hacking and publication bias.

has been declining over time (Brodeur et al. 2020). A solution to this is for both null and positive data to be published in RCT databases going forward, and for journals to institute more rigorous revise and submit processes.

The second main avenue for future work is developing downstream semantic applications which are based on the ontology. One of the simpler applications made possible by this could be an online network of RCTs, in a similar manner to the work of Tang et al. (2008) with Arnetminer, a researcher and research network built on the Friend of a Friend (FOAF) ontology, with users able to search and be recommended RCTs of similar types, whether that be methodology, location or output type. Given that RCT treatment effect data can be successfully gathered from disparate sources as above, advanced downstream applications might include those that can pool treatment effects of specific types of interventions. This would allow for meta-analysis to be performed at a far higher rate than is currently feasible, allowing for both policymakers and researchers to determine effective interventions more efficiently.

Finally, as with the work of Himmelstein (2017), Lui et al. (2020) and Malas et al. (2019), the knowledge graph could potentially be used for downstream research involving machine learning. As in the case of Malas et al. (2019) or Himmelstein (2020), this might involve a classification model that could suggest intervention types for a given context. Alternatively, as in Lui et al. (2020), vectorisation of predicates, objects and subjects using a utility such as graph2vec could provide downstream analysis in identifying symptomatic issues in international development and policy, clustering issues in this domain and allowing policymakers and researchers to identify and define them further.

## 6.4 Final Remarks

The research of this dissertation highlights the need for further integration of data in the social sciences and international development for the purpose of systematic review. The ultimate goal of using these methods is to discover better ways of helping people, be that more effective ways of lifting people out of poverty or nudging individuals towards better health or education outcomes. If policymakers and the international community are to learn the lessons that this research has to offer, the data that RCTs produce must be compiled in a machine readable and replicable way so that conclusions can be neatly and concisely drawn. As the pool of knowledge grows ever larger, it becomes harder to synthesise these lessons without the assistance of well organised data that can be processed by digital means.

# Appendix

## A.1 Ontology Interview Questions

1. Imagine you are about to read a paper on a recently conducted randomised controlled trial (RCT). As generally as possible, what are the key questions that you feel this paper should answer?

   a. What methodological and quality aspects should the research should make clear?

   b. What units of measurement might you expect to be used?

2. What are the key outcomes or results that you look for in a randomised controlled trial?

   a. Do you often come across RCTs with unexpected secondary results?

   b. What types of populations do your RCTs tend to target?

3. Speaking about your field of interest more generally, what do you feel are:

   a. The most pressing questions in the literature that should steer future research.

   b. What are the current lessons that the empirical literature has taught us about your field?

4. What are the main types of methodological and theoretical approaches used for your field?

   a. What are the main methodological approaches you use?

   b. Are there any academic papers on best practices that you would recommend on methodological approaches to RCTs?

   c. Are there any key academics in the area of empirical methodology in your field that I should be aware of?

## A.2 Ethical Approval



((MYzone))

**Matt Murtagh <mmurtagh@tcd.ie>**

**TCD REC WebApp: The status of 'OWL Ontology for Randomised Controlled Trials in Development Economics' (880) has been updated by the Committee**
1 message

**rec-app-help@tchpc.tcd.ie** <rec-app-help@tchpc.tcd.ie> 26 April 2021 at 09:34
To: mmurtagh@tcd.ie

The status of 'OWL Ontology for Randomised Controlled Trials in Development Economics' has been updated by the Committee.

Title: 'OWL Ontology for Randomised Controlled Trials in Development Economics'
Applicant Name: Matt Murtagh-White
Submitted by: Matt Murtagh-White
Academic Supervisor: Declan O'Sullivan
Application Number: 20210103

Result of the REC Meeting: Approved

The Feedback from the Committee is as follows:
This project is now approved, we wish you the best with your study.

The application can be viewed here:

https://webhost.tchpc.tcd.ie/research_ethics/?q=node/880

If amendments are required, please use the following link to edit the application and upload the changes:

https://webhost.tchpc.tcd.ie/research_ethics/?q=node/880/edit

## A.3 Full Ontology Diagram

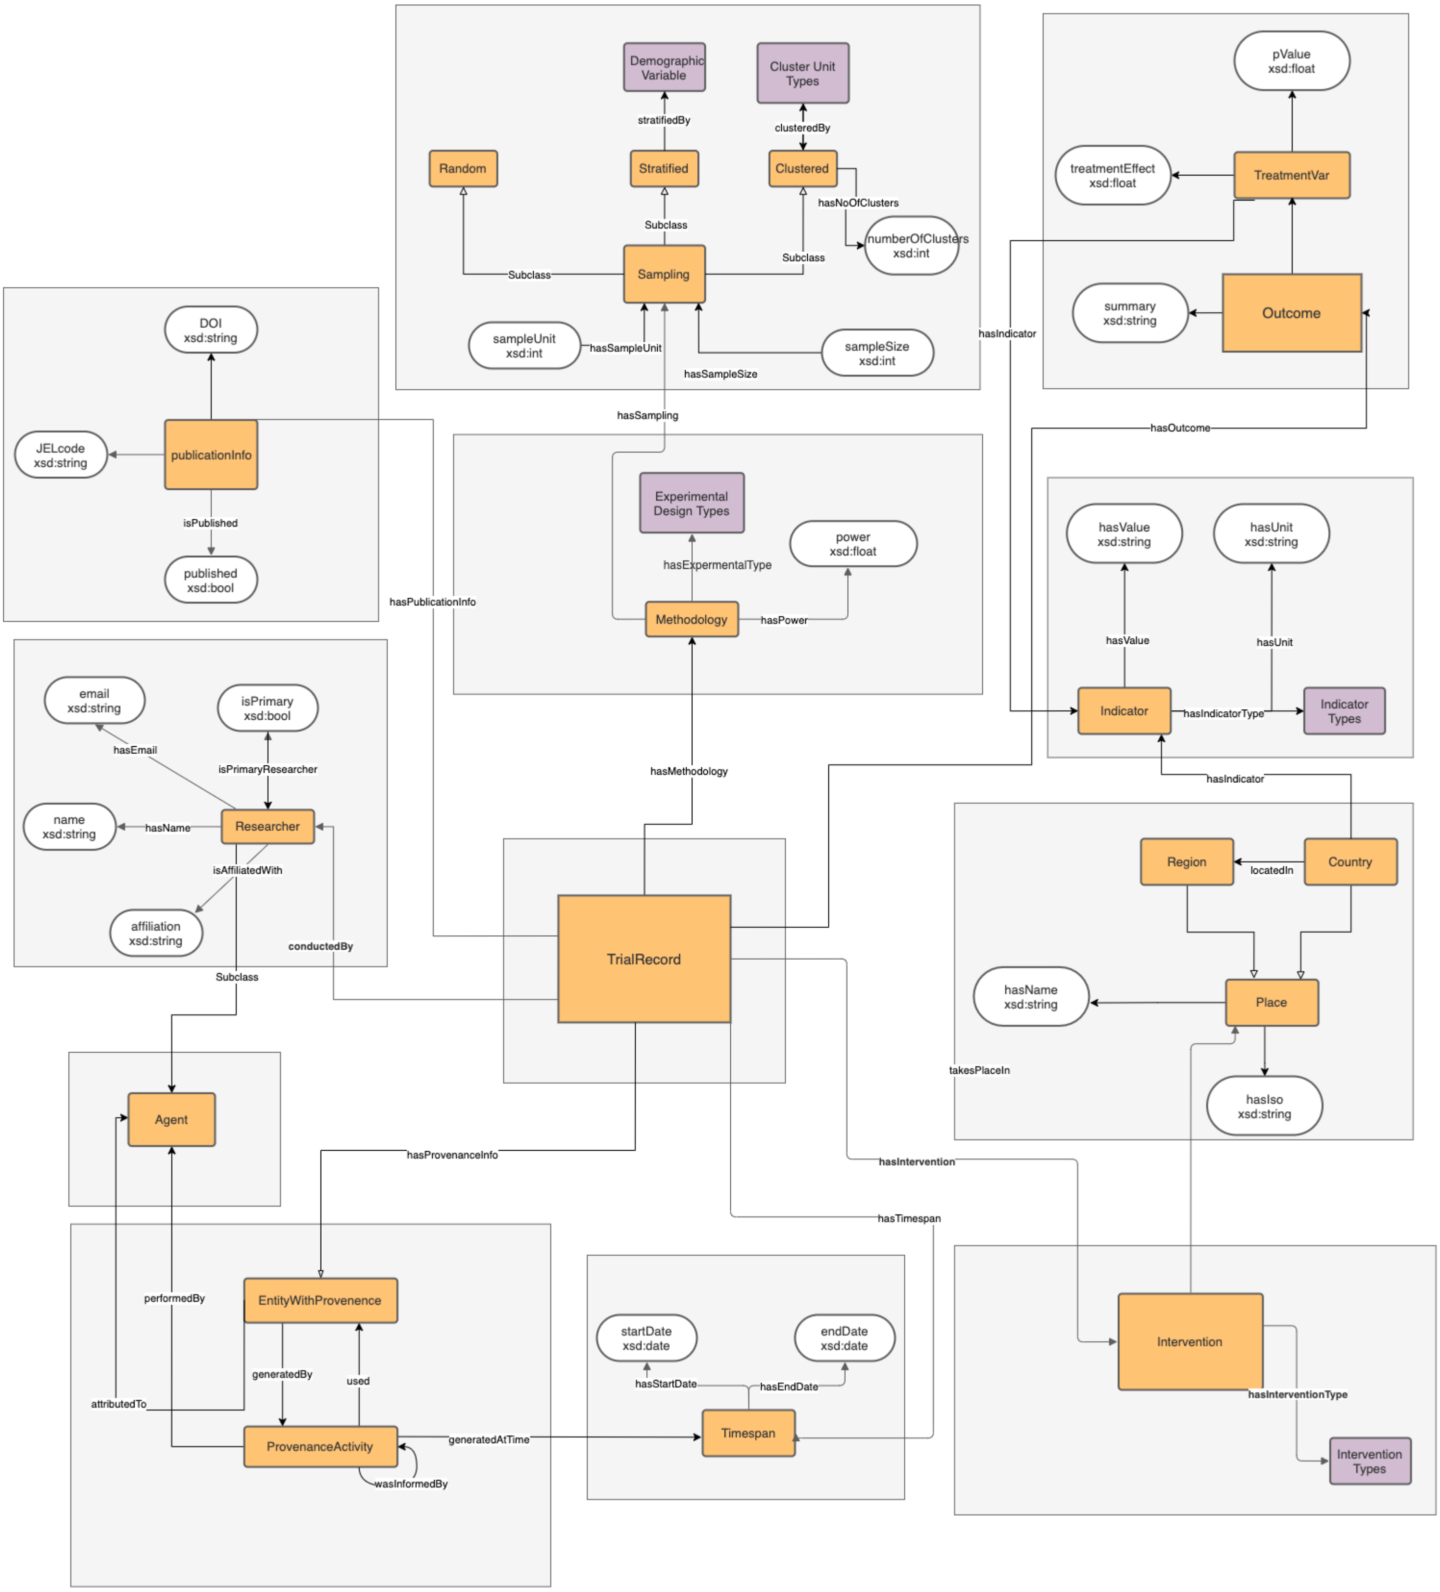

## A.4 Link to Github Repository

https://github.com/mmurtaghw/erct

## A.5 SQARQL Queries for Competency Questions

**CQ1**

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>

PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
select  ?country ?countryCode ?DemVar ?countryName where {
   ?country erct:hasISOAlpha2 ?countryCode.
   ?country erct:hasName ?countryName.
   FILTER (?countryCode IN ("RW", "KE", "UG", "TZ", "BU", "SS", "SD") ).
   ?Stratified a erct:Stratified;
      erct:stratifiedBy ?DemVar.
      FILTER regex(?DemVar, "gender", "i").
}
```

**CQ2**

**Query a:**

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

select  ?unique (COUNT(?treatment) * 0.05 as ?count) where {
         ?treatment erct:takesPlaceIn ?country.
         ?treatment erct:hasInterventionType ?treatmentType.
      ?country erct:hasIndicator ?indicator.
      ?indicator erct:hasValue ?value.
      ?indicator erct:hasIndicatorType ?type.
      ?country erct:hasName ?countryName.
      FILTER (?type = "Net Savings" && ?treatmentType = "finance" && ?value > -10000).
                  }
         GROUP BY ?unique
```

**Query b:**

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

select ?treatment ?country ?countryName ?indicator ?value ?type ?treatmentType where {
         ?treatment erct:takesPlaceIn ?country.
         ?treatment erct:hasInterventionType ?treatmentType.
      ?country erct:hasIndicator ?indicator.
      ?indicator erct:hasValue ?value.
```

```
    ?indicator erct:hasIndicatorType ?type.
    ?country erct:hasName ?countryName.
    FILTER (?type = "Net Savings" && ?treatmentType = "finance" && ?value > -10000).
             }
  ORDER BY DESC(?value) limit 58
```

**CQ3**

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

select ?treatment ?country ?countryName ?indicator ?value ?type ?treatmentType where {
       ?treatment erct:takesPlaceIn ?country.
       ?treatment erct:hasInterventionType ?treatmentType.
    ?country erct:hasIndicator ?indicator.
    ?indicator erct:hasValue ?value.
    ?indicator erct:hasIndicatorType ?type.
    ?country erct:hasName ?countryName.
    FILTER (?type = "HDI" && ?treatmentType = "education"  && ?value > 0.8).
             }
  ORDER BY (?value)
```

**CQ4**

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

select ?treatment ?country ?countryName ?indicator ?value ?type ?treatmentType where {
       ?treatment erct:takesPlaceIn ?country.
       ?treatment erct:hasInterventionType ?treatmentType.
    ?country erct:hasIndicator ?indicator.
    ?indicator erct:hasValue ?value.
    ?indicator erct:hasIndicatorType ?type.
    ?country erct:hasName ?countryName.
    FILTER (?type = "HDI" && ?value < 0.5).
    FILTER regex(?treatmentType, "firms_and_productivity", "i").
             }
  ORDER BY (?value)
```

**CQ5**

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

select ?treatment ?country ?countryName ?indicator ?outcome ?value ?type ?treatmentType
```

```
where {
        ?treatment erct:takesPlaceIn ?country.
        ?treatment erct:hasInterventionType ?treatmentType.
        ?treatment erct:hasOutcome ?outcome.
    ?country erct:hasIndicator ?indicator.
    ?indicator erct:hasValue ?value.
    ?indicator erct:hasIndicatorType ?type.
    ?country erct:hasName ?countryName.
    FILTER (?type = "Enrollment in Primary Education" && ?treatmentType = "education"  &&
?value < 70).
            }
  ORDER BY (?value)
```

**CQ6**

**Query a:**

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

select  ?unique (COUNT(?treatment) * 0.1 as ?count) where {
    ?country erct:hasIndicator ?indicator.
    ?indicator erct:hasValue ?value.
        ?treatment erct:takesPlaceIn ?country.
        ?treatment erct:hasInterventionType ?treatmentType.
    ?indicator erct:hasIndicatorType ?type.
    ?country erct:hasName ?countryName.
    FILTER (?type = "Doctors per 1000").
    FILTER regex(?treatmentType, "HEALTH", "i").
            }
        GROUP BY ?unique
```

**Query b:**

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

select ?treatment ?country ?countryName ?indicator ?value ?type ?treatmentType where {
    ?country erct:hasIndicator ?indicator.
    ?indicator erct:hasValue ?value.
        ?treatment erct:takesPlaceIn ?country.
        ?treatment erct:hasInterventionType ?treatmentType.
    ?indicator erct:hasIndicatorType ?type.
    ?country erct:hasName ?countryName.
    FILTER (?type = "Doctors per 1000").
```

```
    FILTER regex(?treatmentType, "HEALTH", "i").
                  }
    ORDER BY (?value) LIMIT 135
```

**CQ7**

Query a:

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

select  ?unique (COUNT(?treatment)*0.1 as ?count) where {
        ?treatment erct:takesPlaceIn ?country.
        ?treatment erct:hasInterventionType ?treatmentType.
        ?treatment erct:hasOutcome ?outcome.
    ?country erct:hasIndicator ?indicator.
    ?indicator erct:hasValue ?value.
    ?indicator erct:hasIndicatorType ?type.
    ?country erct:hasName ?countryName.
    FILTER (?type = "GDP").
    FILTER regex(?treatmentType, "finance", "i").
                  }
        GROUP BY ?unique
```

**Query b:**

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

select ?treatment ?country ?countryName ?outcome ?indicator ?value ?type ?treatmentType
where {
        ?treatment erct:takesPlaceIn ?country.
        ?treatment erct:hasInterventionType ?treatmentType.
        ?treatment erct:hasOutcome ?outcome.
    ?country erct:hasIndicator ?indicator.
    ?indicator erct:hasValue ?value.
    ?indicator erct:hasIndicatorType ?type.
    ?country erct:hasName ?countryName.
    FILTER (?type = "GDP").
    FILTER regex(?treatmentType, "finance", "i").
                  }
  ORDER BY (?value) limit 108
```

**CQ8**

```
PREFIX erct: <http://www.semanticweb.org/ERCT/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>

select ?treatment ?country ?outcome ?treatmentType where {
          ?treatment erct:takesPlaceIn ?country.
          ?treatment erct:hasInterventionType ?treatmentType.
          ?treatment erct:hasOutcome ?outcome.
      FILTER regex(?treatmentType, "agriculture", "i").
                   }
   ORDER BY (?value) limit 100
```

## A.6 Ontology Expressed as OWL

```
@prefix : <http://www.semanticweb.org/matt/ontologies/2021/1/untitled-ontology-2#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix xml: <http://www.w3.org/XML/1998/namespace> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@base <http://www.semanticweb.org/ERCT> .

<http://www.semanticweb.org/ERCT> rdf:type owl:Ontology .

#################################################################
#    Object Properties
#################################################################

###  http://www.semanticweb.org/ERCT#attributedTo
:attributedTo rdf:type owl:ObjectProperty ;
            rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#clusteredBy
:clusteredBy rdf:type owl:ObjectProperty ;
           rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#conductedBy
:conductedBy rdf:type owl:ObjectProperty ;
           rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#generatedAtTime
:generatedAtTime rdf:type owl:ObjectProperty ;
               rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#generatedBy
:generatedBy rdf:type owl:ObjectProperty ;
           rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#hasExperimentalDesignType
:hasExperimentalDesignType rdf:type owl:ObjectProperty ;
                         rdfs:subPropertyOf owl:topObjectProperty .
```

```
### http://www.semanticweb.org/ERCT#hasIndicator
:hasIndicator rdf:type owl:ObjectProperty ;
        rdfs:subPropertyOf owl:topObjectProperty .


### http://www.semanticweb.org/ERCT#hasIndicatorType
:hasIndicatorType rdf:type owl:ObjectProperty ;
        rdfs:subPropertyOf owl:topObjectProperty .


### http://www.semanticweb.org/ERCT#hasIntervention
:hasIntervention rdf:type owl:ObjectProperty ;
        rdfs:subPropertyOf owl:topObjectProperty .


### http://www.semanticweb.org/ERCT#hasInterventionType
:hasInterventionType rdf:type owl:ObjectProperty ;
        rdfs:subPropertyOf owl:topObjectProperty .


### http://www.semanticweb.org/ERCT#hasMethodology
:hasMethodology rdf:type owl:ObjectProperty ;
        rdfs:subPropertyOf owl:topObjectProperty .


### http://www.semanticweb.org/ERCT#hasOutcome
:hasOutcome rdf:type owl:ObjectProperty ;
        rdfs:subPropertyOf owl:topObjectProperty .


### http://www.semanticweb.org/ERCT#hasProvenenceInfo
:hasProvenenceInfo rdf:type owl:ObjectProperty ;
        rdfs:subPropertyOf owl:topObjectProperty .


### http://www.semanticweb.org/ERCT#hasPublicationInfo
:hasPublicationInfo rdf:type owl:ObjectProperty ;
        rdfs:subPropertyOf owl:topObjectProperty .


### http://www.semanticweb.org/ERCT#hasSampling
:hasSampling rdf:type owl:ObjectProperty ;
        rdfs:subPropertyOf owl:topObjectProperty .


### http://www.semanticweb.org/ERCT#hasTimespan
:hasTimespan rdf:type owl:ObjectProperty ;
```

```
          rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#hasTreatmentVar
:hasTreatmentVar rdf:type owl:ObjectProperty ;
            rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#locatedIn
:locatedIn rdf:type owl:ObjectProperty ;
        rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#performedBy
:performedBy rdf:type owl:ObjectProperty ;
          rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#statifiedBy
:statifiedBy rdf:type owl:ObjectProperty ;
          rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#tookPlaceIn
:tookPlaceIn rdf:type owl:ObjectProperty ;
          rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#used
:used rdf:type owl:ObjectProperty ;
    rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.semanticweb.org/ERCT#wasInformedBy
:wasInformedBy rdf:type owl:ObjectProperty ;
           rdfs:subPropertyOf owl:topObjectProperty .


###  http://www.w3.org/2002/07/owl#topObjectProperty
owl:topObjectProperty rdf:type owl:ObjectProperty ,
                       owl:TransitiveProperty .


#################################################################
#    Data properties
#################################################################
```

```
###  http://www.semanticweb.org/ERCT#hasAffiliation
:hasAffiliation rdf:type owl:DatatypeProperty ;
            rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasControlGroupSize
:hasControlGroupSize rdf:type owl:DatatypeProperty ;
                rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasDOI
:hasDOI rdf:type owl:DatatypeProperty ;
      rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasEmail
:hasEmail rdf:type owl:DatatypeProperty ;
        rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasEndDate
:hasEndDate rdf:type owl:DatatypeProperty ;
         rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasISO
:hasISO rdf:type owl:DatatypeProperty ;
      rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasJELcode
:hasJELcode rdf:type owl:DatatypeProperty ;
         rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasName
:hasName rdf:type owl:DatatypeProperty ;
       rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasNoOfClusters
:hasNoOfClusters rdf:type owl:DatatypeProperty ;
             rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasPValue
:hasPValue rdf:type owl:DatatypeProperty ;
```

```
        rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasPower
:hasPower rdf:type owl:DatatypeProperty ;
        rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasSampleSize
:hasSampleSize rdf:type owl:DatatypeProperty ;
            rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasStartDate
:hasStartDate rdf:type owl:DatatypeProperty ;
           rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasTreatmentEffect
:hasTreatmentEffect rdf:type owl:DatatypeProperty ;
                rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasTreatmentGroupSize
:hasTreatmentGroupSize rdf:type owl:DatatypeProperty ;
                  rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasUnit
:hasUnit rdf:type owl:DatatypeProperty ;
       rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#hasValue
:hasValue rdf:type owl:DatatypeProperty ;
        rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#isControl
:isControl rdf:type owl:DatatypeProperty ;
         rdfs:subPropertyOf owl:topDataProperty .


###  http://www.semanticweb.org/ERCT#isPrimaryREsearcher
:isPrimaryREsearcher rdf:type owl:DatatypeProperty ;
                rdfs:subPropertyOf owl:topDataProperty .
```

```
###  http://www.semanticweb.org/ERCT#isPublished
:isPublished rdf:type owl:DatatypeProperty ;
         rdfs:subPropertyOf owl:topDataProperty .


#################################################################
#    Classes
#################################################################

###  http://www.semanticweb.org/ERCT#Agent
:Agent rdf:type owl:Class .


###  http://www.semanticweb.org/ERCT#ClusterUnitType
:ClusterUnitType rdf:type owl:Class .


###  http://www.semanticweb.org/ERCT#Clustered
:Clustered rdf:type owl:Class ;
        rdfs:subClassOf :Sampling ,
                    [ rdf:type owl:Restriction ;
                      owl:onProperty :clusteredBy ;
                      owl:allValuesFrom :ClusterUnitType
                    ] ,
                    [ rdf:type owl:Restriction ;
                      owl:onProperty :hasNoOfClusters ;
                      owl:allValuesFrom xsd:integer
                    ] .


###  http://www.semanticweb.org/ERCT#Country
:Country rdf:type owl:Class ;
       owl:equivalentClass [ rdf:type owl:Restriction ;
                     owl:onProperty :hasIndicator ;
                     owl:someValuesFrom :Indicator
                    ] ;
       rdfs:subClassOf :Place .


###  http://www.semanticweb.org/ERCT#EntityWithProvenence
:EntityWithProvenence rdf:type owl:Class ;
              rdfs:subClassOf [ rdf:type owl:Restriction ;
                        owl:onProperty :attributedTo ;
                        owl:someValuesFrom :Agent
                      ] ,
                      [ rdf:type owl:Restriction ;
```

```
                                   owl:onProperty :generatedAtTime ;
                                   owl:allValuesFrom :Timespan
                                 ] ,
                                 [ rdf:type owl:Restriction ;
                                   owl:onProperty :generatedBy ;
                                   owl:allValuesFrom :ProvenanceActivity
                                 ] .


###  http://www.semanticweb.org/ERCT#ExperimentalDesignType
:ExperimentalDesignType rdf:type owl:Class .


###  http://www.semanticweb.org/ERCT#Indicator
:Indicator rdf:type owl:Class ;
           rdfs:subClassOf [ rdf:type owl:Restriction ;
                             owl:onProperty :hasIndicatorType ;
                             owl:qualifiedCardinality "1"^^xsd:nonNegativeInteger ;
                             owl:onClass :IndicatorTypes
                           ] ,
                           [ rdf:type owl:Restriction ;
                             owl:onProperty :hasUnit ;
                             owl:qualifiedCardinality "1"^^xsd:nonNegativeInteger ;
                             owl:onDataRange xsd:string
                           ] ,
                           [ rdf:type owl:Restriction ;
                             owl:onProperty :hasValue ;
                             owl:qualifiedCardinality "1"^^xsd:nonNegativeInteger ;
                             owl:onDataRange xsd:float
                           ] .


###  http://www.semanticweb.org/ERCT#IndicatorTypes
:IndicatorTypes rdf:type owl:Class .


###  http://www.semanticweb.org/ERCT#Intervention
:Intervention rdf:type owl:Class ;
              rdfs:subClassOf [ rdf:type owl:Restriction ;
                                owl:onProperty :hasIntervention ;
                                owl:someValuesFrom :Intervention
                              ] ,
                              [ rdf:type owl:Restriction ;
                                owl:onProperty :tookPlaceIn ;
                                owl:someValuesFrom :Country
                              ] ,
                              [ rdf:type owl:Restriction ;
```

```
                    owl:onProperty :hasIndicator ;
                    owl:allValuesFrom :TreatmentVar
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasTimespan ;
                    owl:allValuesFrom :Timespan
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasInterventionType ;
                    owl:minQualifiedCardinality "1"^^xsd:nonNegativeInteger ;
                    owl:onClass :InterventionTypes
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasSampleSize ;
                    owl:allValuesFrom xsd:int
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :isControl ;
                    owl:allValuesFrom xsd:boolean
                  ] .


###  http://www.semanticweb.org/ERCT#InterventionTypes
:InterventionTypes rdf:type owl:Class .


###  http://www.semanticweb.org/ERCT#Methodology
:Methodology rdf:type owl:Class ;
        rdfs:subClassOf [ rdf:type owl:Restriction ;
                    owl:onProperty :hasExperimentalDesignType ;
                    owl:someValuesFrom :ExperimentalDesignType
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasSampling ;
                    owl:allValuesFrom :Sampling
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasPower ;
                    owl:someValuesFrom xsd:float
                  ] .


###  http://www.semanticweb.org/ERCT#Outcome
:Outcome rdf:type owl:Class ;
      rdfs:subClassOf [ rdf:type owl:Restriction ;
                  owl:onProperty :hasIndicator ;
                  owl:someValuesFrom :TreatmentVar
```

```
                ] .


### http://www.semanticweb.org/ERCT#Place
:Place rdf:type owl:Class ;
    rdfs:subClassOf [ rdf:type owl:Restriction ;
                owl:onProperty :hasISO ;
                owl:qualifiedCardinality "1"^^xsd:nonNegativeInteger ;
                owl:onDataRange xsd:string
               ] ,
               [ rdf:type owl:Restriction ;
                owl:onProperty :hasName ;
                owl:qualifiedCardinality "1"^^xsd:nonNegativeInteger ;
                owl:onDataRange xsd:string
               ] .


### http://www.semanticweb.org/ERCT#ProvenanceActivity
:ProvenanceActivity rdf:type owl:Class ;
            rdfs:subClassOf [ rdf:type owl:Restriction ;
                        owl:onProperty :wasInformedBy ;
                        owl:someValuesFrom :ProvenanceActivity
                       ] ,
                       [ rdf:type owl:Restriction ;
                        owl:onProperty :used ;
                        owl:allValuesFrom :EntityWithProvenence
                       ] .


### http://www.semanticweb.org/ERCT#Random
:Random rdf:type owl:Class ;
     rdfs:subClassOf :Sampling .


### http://www.semanticweb.org/ERCT#Region
:Region rdf:type owl:Class ;
     rdfs:subClassOf :Place ,
               [ rdf:type owl:Restriction ;
                owl:onProperty :locatedIn ;
                owl:qualifiedCardinality "1"^^xsd:nonNegativeInteger ;
                owl:onClass :Country
               ] .


### http://www.semanticweb.org/ERCT#Researcher
:Researcher rdf:type owl:Class ;
       rdfs:subClassOf :Agent ,
```

```
                [ rdf:type owl:Restriction ;
                  owl:onProperty :hasAffiliation ;
                  owl:someValuesFrom xsd:string
                ] ,
                [ rdf:type owl:Restriction ;
                  owl:onProperty :hasEmail ;
                  owl:someValuesFrom xsd:boolean
                ] ,
                [ rdf:type owl:Restriction ;
                  owl:onProperty :hasEmail ;
                  owl:someValuesFrom xsd:string
                ] ,
                [ rdf:type owl:Restriction ;
                  owl:onProperty :hasName ;
                  owl:allValuesFrom xsd:string
                ] .


###  http://www.semanticweb.org/ERCT#Sampling
:Sampling rdf:type owl:Class ;
      rdfs:subClassOf [ rdf:type owl:Restriction ;
                owl:onProperty :hasSampleSize ;
                owl:qualifiedCardinality "1"^^xsd:nonNegativeInteger ;
                owl:onDataRange xsd:int
               ] ,
               [ rdf:type owl:Restriction ;
                owl:onProperty :hasUnit ;
                owl:qualifiedCardinality "1"^^xsd:nonNegativeInteger ;
                owl:onDataRange xsd:int
               ] .


###  http://www.semanticweb.org/ERCT#Stratified
:Stratified rdf:type owl:Class ;
       rdfs:subClassOf :Sampling ,
                [ rdf:type owl:Restriction ;
                  owl:onProperty :statifiedBy ;
                  owl:allValuesFrom :demographicVariable
                ] .


###  http://www.semanticweb.org/ERCT#Timespan
:Timespan rdf:type owl:Class ;
      rdfs:subClassOf [ rdf:type owl:Restriction ;
                owl:onProperty :hasEndDate ;
                owl:someValuesFrom xsd:dateTime
               ] ,
```

```
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasStartDate ;
                    owl:allValuesFrom xsd:dateTime
                  ] .


###  http://www.semanticweb.org/ERCT#TreatmentVar
:TreatmentVar rdf:type owl:Class ;
          rdfs:subClassOf [ rdf:type owl:Restriction ;
                    owl:onProperty :hasIndicator ;
                    owl:qualifiedCardinality "1"^^xsd:nonNegativeInteger ;
                    owl:onClass :Indicator
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasPValue ;
                    owl:someValuesFrom xsd:float
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasTreatmentEffect ;
                    owl:someValuesFrom xsd:float
                  ] .


###  http://www.semanticweb.org/ERCT#TrialRecord
:TrialRecord rdf:type owl:Class ;
          rdfs:subClassOf [ rdf:type owl:Restriction ;
                    owl:onProperty :hasOutcome ;
                    owl:someValuesFrom :Outcome
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :conductedBy ;
                    owl:allValuesFrom :Researcher
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasIntervention ;
                    owl:allValuesFrom :Intervention
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasProvenenceInfo ;
                    owl:allValuesFrom :EntityWithProvenence
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasPublicationInfo ;
                    owl:allValuesFrom :publicationInfo
                  ] ,
                  [ rdf:type owl:Restriction ;
                    owl:onProperty :hasSampling ;
```

```
                  owl:allValuesFrom :Sampling
                ] .


###  http://www.semanticweb.org/ERCT#demographicVariable
:demographicVariable rdf:type owl:Class .


###  http://www.semanticweb.org/ERCT#publicationInfo
:publicationInfo rdf:type owl:Class ;
          rdfs:subClassOf [ rdf:type owl:Restriction ;
                   owl:onProperty :hasDOI ;
                   owl:someValuesFrom xsd:string
                 ] ,
                 [ rdf:type owl:Restriction ;
                   owl:onProperty :hasJELcode ;
                   owl:someValuesFrom xsd:string
                 ] ,
                 [ rdf:type owl:Restriction ;
                   owl:onProperty :isPublished ;
                   owl:someValuesFrom xsd:boolean
                 ] .


###  Generated by the OWL API (version 4.5.9.2019-02-01T07:24:44Z) https://github.com/owlcs/owlapi
```

# Bibliography


American Economic Association. (2021a, July 13). Trial Data Access. *Trial Data Access*. https://www.socialscienceregistry.org/site/data

American Economic Association. (2021b, October 7). *JEL Clasification System / EconLit Subject Descriptors*. https://www.aeaweb.org/econlit/jelCodes.php?view=jel

Arenas, M., Bertails, A., Prud'hommeaux, E., & Sequeda, J. (2012). A direct mapping of relational data to RDF. *W3C Recommendation*, *27*, 1–11.

Baird, S., Ferreira, F. H., Özler, B., & Woolcock, M. (2013). Relative effectiveness of conditional and unconditional cash transfers for schooling outcomes in developing countries: A systematic review. *Campbell Systematic Reviews*, *9*(1), 1–124.

Banerjee, A. V., Duflo, E., & Kremer, M. (2016). The influence of randomized controlled trials on development economics research and on development policy. *The State of Economics, The State of the World*, 482–488.

Bastagli, F., Hagen-Zanker, J., Harman, L., Barca, V., Sturge, G., Schmidt, T., & Pellerano, L. (2016). Cash transfers: What does the evidence say. *A Rigorous Review of Programme Impact and the Role of Design and Implementation Features. London: ODI*, *1*(7).

Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The semantic web. *Scientific American*, *284*(5), 34–43.

Bizer, C., Heath, T., & Berners-Lee, T. (2011). Linked data: The story so far. In *Semantic services, interoperability and web applications: Emerging concepts* (pp. 205–227). IGI global.

Bizer, C., Heath, T., & Berners-Lee, T. (2008). *Linked data: Principles and state of the art*. *1*, 40.

Brandt, K. S., & de Boer, V. (2015). Linked data for the international aid transparency initiative. *Journal on Data Semantics*, *4*(3), 187–211.

Brickley, D., Guha, R. V., & McBride, B. (2014). RDF Schema 1.1. *W3C Recommendation*, *25*, 2004–2014.

Brodeur, A., Cook, N., & Heyes, A. (2020). Methods matter: P-hacking and publication bias in causal analysis in economics. *American Economic Review*, *110*(11), 3634–3660.

Cameron, D. B., Mishra, A., & Brown, A. N. (2016). The growth of impact evaluation for international development: How much have we learned? *Journal of Development Effectiveness*, *8*(1), 1–21.

Das, S., Sundara, S., & Cyganiak, R. (2012). R2RML: RDB to RDF mapping language (2012). *URL Http://Www. W3. Org/TR/2012/REC-R2rml-20120927*.

DBpedia. (2021, September 7). *DBpedia Ontology*. https://www.dbpedia.org/resources/ontology/

de Boer, V. (n.d.). *Bringing Knowledge Graphs to the Rural Poor*.

De Boer, V., Gyan, N. B., Bon, A., Tuyp, W., Van Aart, C., & Akkermans, H. (2015). A dialogue with linked data: Voice-based access to market data in the sahel. *Semantic Web*, *6*(1), 23–33.

de Walque, D., Fernald, L., Gertler, P., & Hidrobo, M. (2017). *Cash transfers and child and adolescent development*.

Debruyne, C., & O'Sullivan, D. (2016). R2RML-F: towards sharing and executing domain logic in R2RML mappings. *LDOW@ WWW*.

Duflo, E., & Banerjee, A. (2017). *Handbook of field experiments*. Elsevier.

Duvendack, M., Hombrados, J. G., Palmer-Jones, R., & Waddington, H. (2012). Assessing 'what works' in international development: Meta-analysis for sophisticated dummies. *Journal of Development Effectiveness*, *4*(3), 456–471.

Edwards, R. T., Charles, J. M., & Lloyd-Williams, H. (2013). Public health economics: A systematic review of guidance for the economic evaluation of public health interventions and discussion of key methodological issues. *BMC Public Health*, *13*(1), 1–13.

Ehrlinger, L., & Wöß, W. (2016). Towards a Definition of Knowledge Graphs. *SEMANTiCS (Posters, Demos, SuCCESS)*, *48*(1–4), 2.

El-Sappagh, S., Franda, F., Ali, F., & Kwak, K.-S. (2018). SNOMED CT standard ontology based on the ontology for general medical science. *BMC Medical Informatics and Decision Making*, *18*(1), 1–19.

Fernández-López, M., Poveda-Villalón, M., Suárez-Figueroa, M. C., & Gómez-Pérez, A. (2019). Why are ontologies not reused across the same domain? *Journal of Web Semantics*, *57*, 100492.

Fryer Jr, R. G. (2017). The production of human capital in developed countries: Evidence from 196 randomized field experiments. In *Handbook of economic field experiments* (Vol. 2, pp. 95–322). Elsevier.

Furnham, A., & Boo, H. C. (2011). A literature review of the anchoring effect. *The Journal of Socio-Economics*, *40*(1), 35–42.

Gangemi, A. (2005). *Ontology design patterns for semantic web content*. 262–276.

Gangemi, A., & Presutti, V. (2009). Ontology design patterns. In *Handbook on ontologies* (pp. 221–243). Springer.

Goldbeck, J., & Rothstein, M. (2008). Linking social Networks on the web with FOAF. *Proceedings of the Twenty-Third Conference on Artificial Intelligence, AAA08*.

Grüninger, M., & Fox, M. S. (1995). The role of competency questions in enterprise engineering. In *Benchmarking—Theory and practice* (pp. 22–31). Springer.

Guéret, C., de Boer, V., & Schlobach, S. (2014). Let’s" Downscale" Linked Data. *IEEE Internet Computing*, *18*(2), 70–73.

Hampton, S., & Adams, R. (2018). Behavioural economics vs social practice theory: Perspectives from inside the United Kingdom government. *Energy Research & Social Science*, *46*, 214–224.

Himmelstein, D. S., Lizee, A., Hessler, C., Brueggeman, L., Chen, S. L., Hadley, D., Green, A., Khankhanian, P., & Baranzini, S. E. (2017). Systematic integration of biomedical knowledge prioritizes drugs for repurposing. *Elife*, *6*, e26726.

Hitzler, P., & Krisnadhi, A. (2018). A tutorial on modular ontology modeling with ontology design patterns: The cooking recipes ontology. *ArXiv Preprint ArXiv:1808.08433*.

Holsapple, C. W., & Joshi, K. D. (2002). A collaborative approach to ontology design. *Communications of the ACM*, *45*(2), 42–47.

Hopf, C. (2004). Qualitative interviews: An overview. *A Companion to Qualitative Research*, *203*(8).

Huang, X., Lin, J., & Demner-Fushman, D. (2006). Evaluation of PICO as a knowledge representation for clinical questions. *AMIA Annual Symposium Proceedings*, *2006*, 359.

International Initiative for Impact Evaluation 2021). (2021, July 13). About RIDIE. *About RIDIE*. https://ridie.3ieimpact.org/index.php?r=site/page&view=about

Jack, W., & Suri, T. (2011). *Mobile money: The economics of M-PESA*. National Bureau of Economic Research.

Junior, A. C., Debruyne, C., & O'Sullivan, D. (2016). Incorporating functions in mappings to facilitate the uplift of CSV files into RDF. *European Semantic Web Conference*, 55–59.

Kardas, M., Czapla, P., Stenetorp, P., Ruder, S., Riedel, S., Taylor, R., & Stojnic, R. (2020). Axcell: Automatic extraction of results from machine learning papers. *ArXiv Preprint ArXiv:2004.14356*.

Karima, N., Hammar, K., & Hitzler, P. (2017). How to document ontology design patterns. *Advances in Ontology Design and Patterns*, *32*, 15–28.

Kugler, A. D., & Rojas, I. (2018). *Do CCTs improve employment and earnings in the very long-term? Evidence from Mexico*. National Bureau of Economic Research.

Lei, Z., Sun, Y., Nanehkaran, Y. A., Yang, S., Islam, M. S., Lei, H., & Zhang, D. (2020). A novel data-driven robust framework based on machine learning and knowledge graph for disease classification. *Future Generation Computer Systems*, *102*, 534–548.

Lin, F. P., Groza, T., Kocbek, S., Antezana, E., & Epstein, R. J. (2018). Cancer Care Treatment Outcome Ontology: A Novel Computable Ontology for Profiling Treatment Outcomes in Patients With Solid Tumors. *JCO Clinical Cancer Informatics*, *2*, 1–14.

Liu, Y., Bai, K., Mitra, P., & Giles, C. L. (2007). *Tableseer: Automatic table metadata extraction and searching in digital libraries*. 91–100.

Maki, U., & Mäki, U. (2001). *The economic world view: Studies in the ontology of economics*. Cambridge University Press.

Malas, T. B., Vlietstra, W. J., Kudrin, R., Starikov, S., Charrout, M., Roos, M., Peters, D. J., Kors, J. A., Vos, R., & AC't Hoen, P. (2019). Drug prioritization using the semantic properties of a knowledge graph. *Scientific Reports*, *9*(1), 1–10.

Manola, F., Miller, E., & McBride, B. (2004). RDF primer. *W3C Recommendation*, *10*(1–107), 6.

Mavergames, C., Oliver, S., & Becker, L. (2013). Systematic Reviews as an Interface to the Web of (Trial) Data: Using PICO as an Ontology for Knowledge Synthesis in Evidence-based Healthcare Research. *SePublica*, *994*, 22–26.

McGuinness, D. L., & Van Harmelen, F. (2004). OWL web ontology language overview. *W3C Recommendation*, *10*(10), 2004.

Michie, S., & Johnston, M. (2017). Optimising the value of the evidence generated in implementation science: The use of ontologies to address the challenges. *Implementation Science*, *12*(1), 1–4.

Miguel, E., & Kremer, M. (2004). Worms: Identifying impacts on education and health in the presence of treatment externalities. *Econometrica*, *72*(1), 159–217.

Musen, M. A. (2015). The protégé project: A look back and a look forward. *AI Matters*, *1*(4), 4–12.

National Center for Biomedical Ontology. (2021, August 28). *Bioportal*. https://bioportal.bioontology.org/ontologies

Neubert, J. (2009). Bringing the" Thesaurus for Economics" on to the Web of Linked Data. *LDOW*, *25964*, 102.

Novemsky, N., & Kahneman, D. (2005). The boundaries of loss aversion. *Journal of Marketing Research*, *42*(2), 119–128.

Oliver, A. (2013). From nudging to budging: Using behavioural economics to inform public sector policy. *Journal of Social Policy*, *42*(4), 685–700.

Ontotext. (2021, July 13). GraphDB. *GraphDB*. https://graphdb.ontotext.com/documentation/free/

Parker, S. W., & Vogl, T. (2018). *Do conditional cash transfers improve economic outcomes in the next generation? Evidence from Mexico*. National Bureau of Economic Research.

Paulheim, H. (2017). Knowledge graph refinement: A survey of approaches and evaluation methods. *Semantic Web*, *8*(3), 489–508.

Pernisch, R., Dell'Aglio, D., & Bernstein, A. (2021). Beware of the hierarchy—An analysis of ontology evolution and the materialisation impact for biomedical ontologies. *Journal of Web Semantics*, 100658. https://doi.org/10.1016/j.websem.2021.100658

Peters, J., Langbein, J., & Roberts, G. (2016). Policy evaluation, randomized controlled trials, and external validity—A systematic review. *Economics Letters*, *147*, 51–54.

Presutti, V., Daga, E., Gangemi, A., & Blomqvist, E. (2009). EXtreme design with content ontology design patterns. *Proc. Workshop on Ontology Patterns*, 83–97.

Rawlings, L. B., & Rubio, G. M. (2005). Evaluating the impact of conditional cash transfer programs. *The World Bank Research Observer*, *20*(1), 29–55.

Reisch, L. A., & Zhao, M. I. N. (2017). Behavioural economics, consumer behaviour and consumer policy: State of the art. *Behavioural Public Policy*, *1*(2), 190.

Riley, E. (2018). Mobile money and risk sharing against village shocks. *Journal of Development Economics*, *135*, 43–58.

Ristoski, P., Rosati, J., Di Noia, T., De Leone, R., & Paulheim, H. (2019). RDF2Vec: RDF graph embeddings and their applications. *Semantic Web*, *10*(4), 721–752.

Robinson, O. C. (2014). Sampling in interview-based qualitative research: A theoretical and practical guide. *Qualitative Research in Psychology*, *11*(1), 25–41.

Rowe, M. (2009). *Interlinking distributed social graphs*. LDOW.

Saavedra, J. E., & García, S. (2012). Impacts of conditional cash transfer programs on educational outcomes in developing countries: A meta-analysis. *RAND Labor and Population Working Paper Series, WR-921-1*.

Schardt, C., Adams, M. B., Owens, T., Keitz, S., & Fontelo, P. (2007). Utilization of the PICO framework to improve searching PubMed for clinical questions. *BMC Medical Informatics and Decision Making*, *7*(1), 1–6.

Seaborne, A., Manjunath, G., Bizer, C., Breslin, J., Das, S., Davis, I., Harris, S., Idehen, K., Corby, O., & Kjernsmo, K. (2008). SPARQL/Update: A language for updating RDF graphs. *W3c Member Submission*, *15*.

Seaborne, A., & Prud'hommeaux, E. (2006). SPARQL query language for RDF. *W3C Recommendation (January 2008)*.

Shadbolt, N., & O'Hara, K. (2013). Linked data in government. *IEEE Internet Computing*, *17*(4), 72–77.

Shimizu, C., Hitzler, P., Hirt, Q., Rehberger, D., Estrecha, S. G., Foley, C., Sheill, A. M., Hawthorne, W., Mixter, J., & Watrall, E. (2020). The enslaved ontology: Peoples of the historic slave trade. *Journal of Web Semantics*, *63*, 100567.

Sim, I., Olasov, B., & Carini, S. (2004). An ontology of randomized controlled trials for evidence-based practice: Content specification and evaluation using the competency decomposition method. *Journal of Biomedical Informatics*, *37*(2), 108–119.

Sim, I., Tu, S. W., Carini, S., Lehmann, H. P., Pollock, B. H., Peleg, M., & Wittkowski, K. M. (2014). The Ontology of Clinical Research (OCRe): An informatics foundation for the science of clinical research. *Journal of Biomedical Informatics*, *52*, 78–91.

Sukhtankar, S. (2017). Replications in development economics. *American Economic Review*, *107*(5), 32–36.

Tang, J., Zhang, J., Yao, L., Li, J., Zhang, L., & Su, Z. (2008). *Arnetminer: Extraction and mining of academic social networks*. 990–998.

Thaler, R., & Sunstein, C. (2009). NUDGE: Improving Decisions About Health, Wealth, and Happiness. In *Nudge: Improving Decisions about Health, Wealth, and Happiness* (Vol. 47).

The World Bank, T. W. B. (2021, July 13). World Bank Data API. *World Bank Data API*. https://datahelpdesk.worldbank.org/knowledgebase/topics/125589-developer-information

Tudorache, T. (2020). Ontology engineering: Current state, challenges, and future directions. *Semantic Web*, *11*(1), 125–138.

United Nations. (2020, January 1). Knowledge Graphs for Social Good. *Knowledge Graphs for Social Good*. https://knowledgegraphsocialgood.pubpub.org/schedule

Valkering, O., de Boer, V., Lô, G., Blankendaal, R., & Schlobach, S. (2016). *The semantic web in an sms*. 697–712.

Waddington, H., White, H., Snilstveit, B., Hombrados, J. G., Vojtkova, M., Davies, P., Bhavsar, A., Eyers, J., Koehlmoos, T. P., & Petticrew, M. (2012). How to do a good systematic review of effects in international development: A tool kit. *Journal of Development Effectiveness*, *4*(3), 359–387.

Williams, K., Wu, J., Choudhury, S. R., Khabsa, M., & Giles, C. L. (2014). *Scholarly big data information extraction and integration in the citeseer χ digital library*. 68–73.

World Health Organization. (2016). *Annual technical report 2015: Department of reproductive health and research, including UNDP/UNFPA/WHO/World Bank Special Programme of Research Training in Human Reproduction (HRP)*. World Health Organization.